Extreme imaging using cell phones

ICCV Extreme Imaging Workshop
December 17, 2015
(revised September 15, 2016)

Marc Levoy
Principal Engineer
Google Research

Professor, Emeritus
Computer Science Department
Stanford University

Really? Extreme imaging using cell phones?

- small aperture
  - limits the number of photons it can gather
  - noisy images in low light
- small sensor
  - limits the number of electrons each pixel can store
  - poor dynamic range
- small lens module
  - limits the length of its optical path
  - no telephoto, no zoom

Jesse Levinson Canon 10D, 28mm f/4, 3 min, ISO 100, 4 image pano

(video available on YouTube at https://youtu.be/S7lbnMd56Ys, but the app is not currently available to the public)

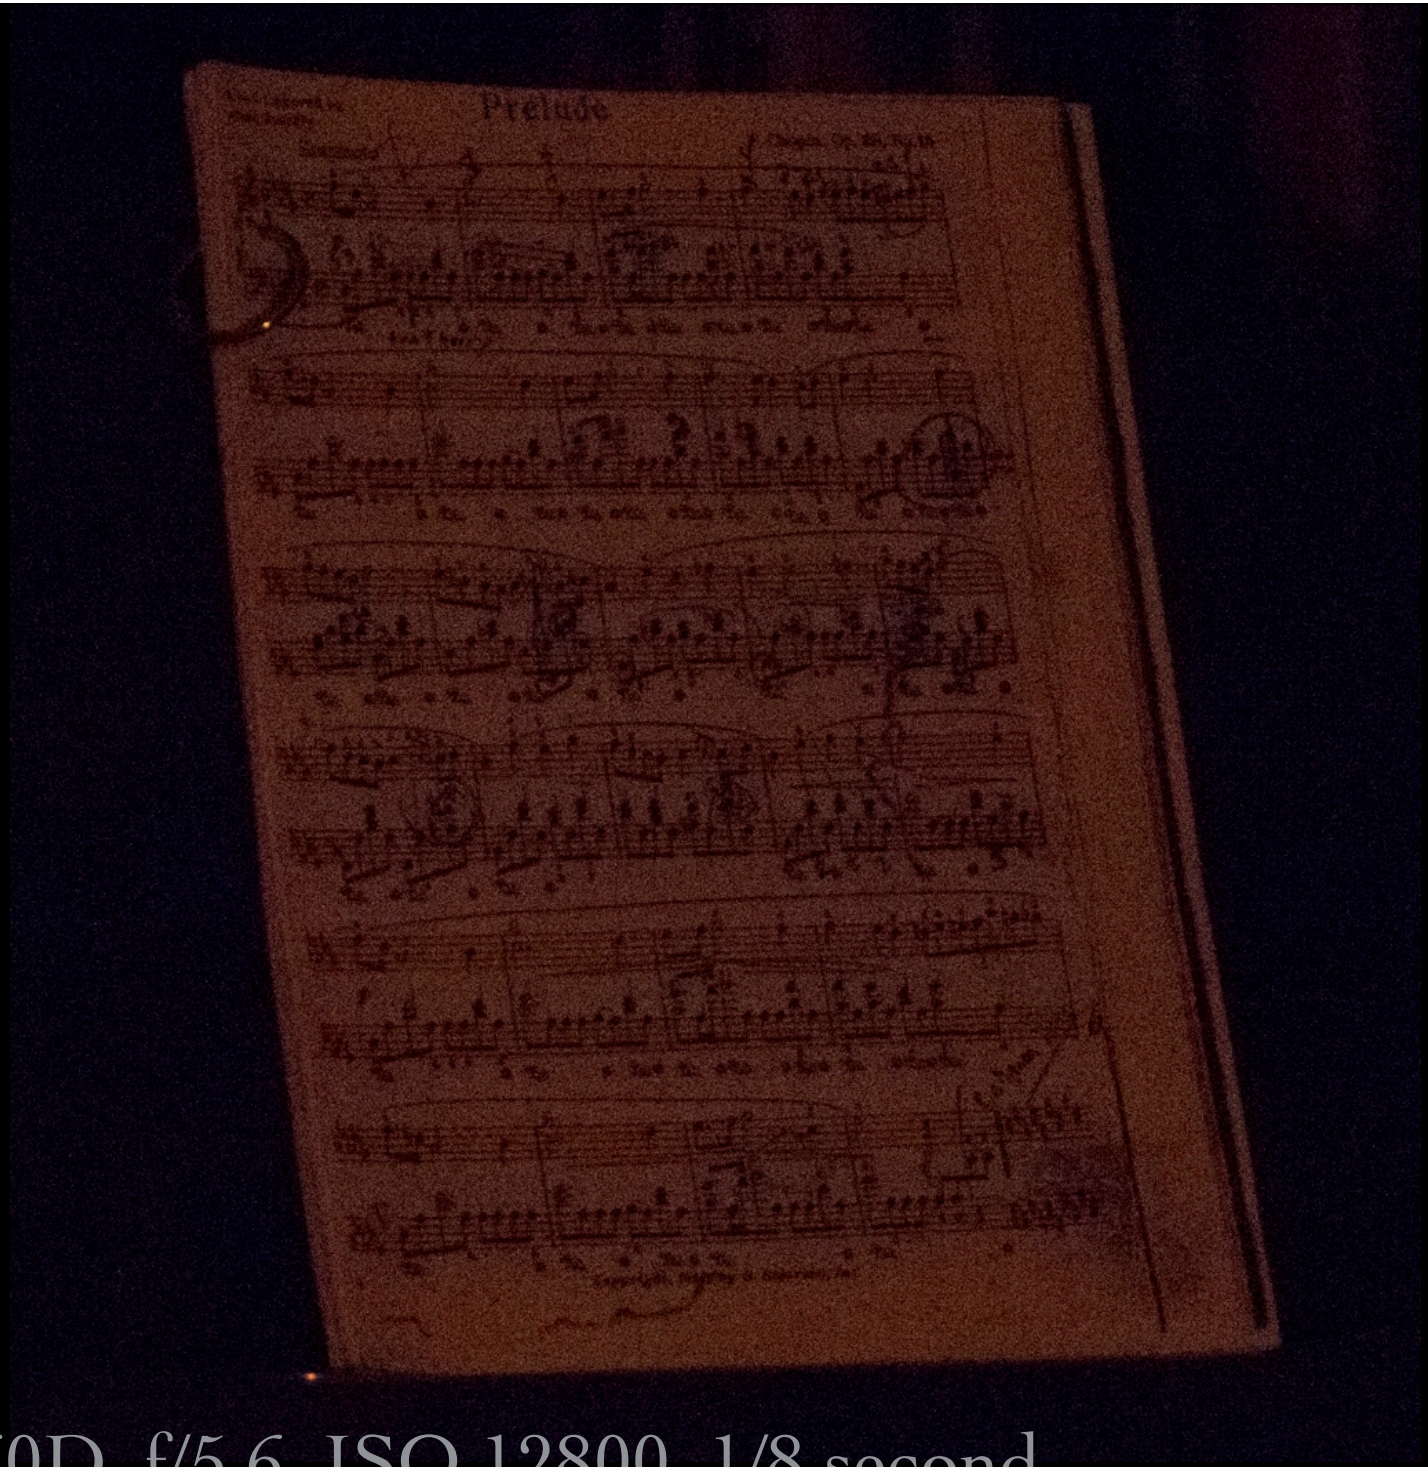

Canon 60D, f/5.6, ISO 12800, 1/8 second

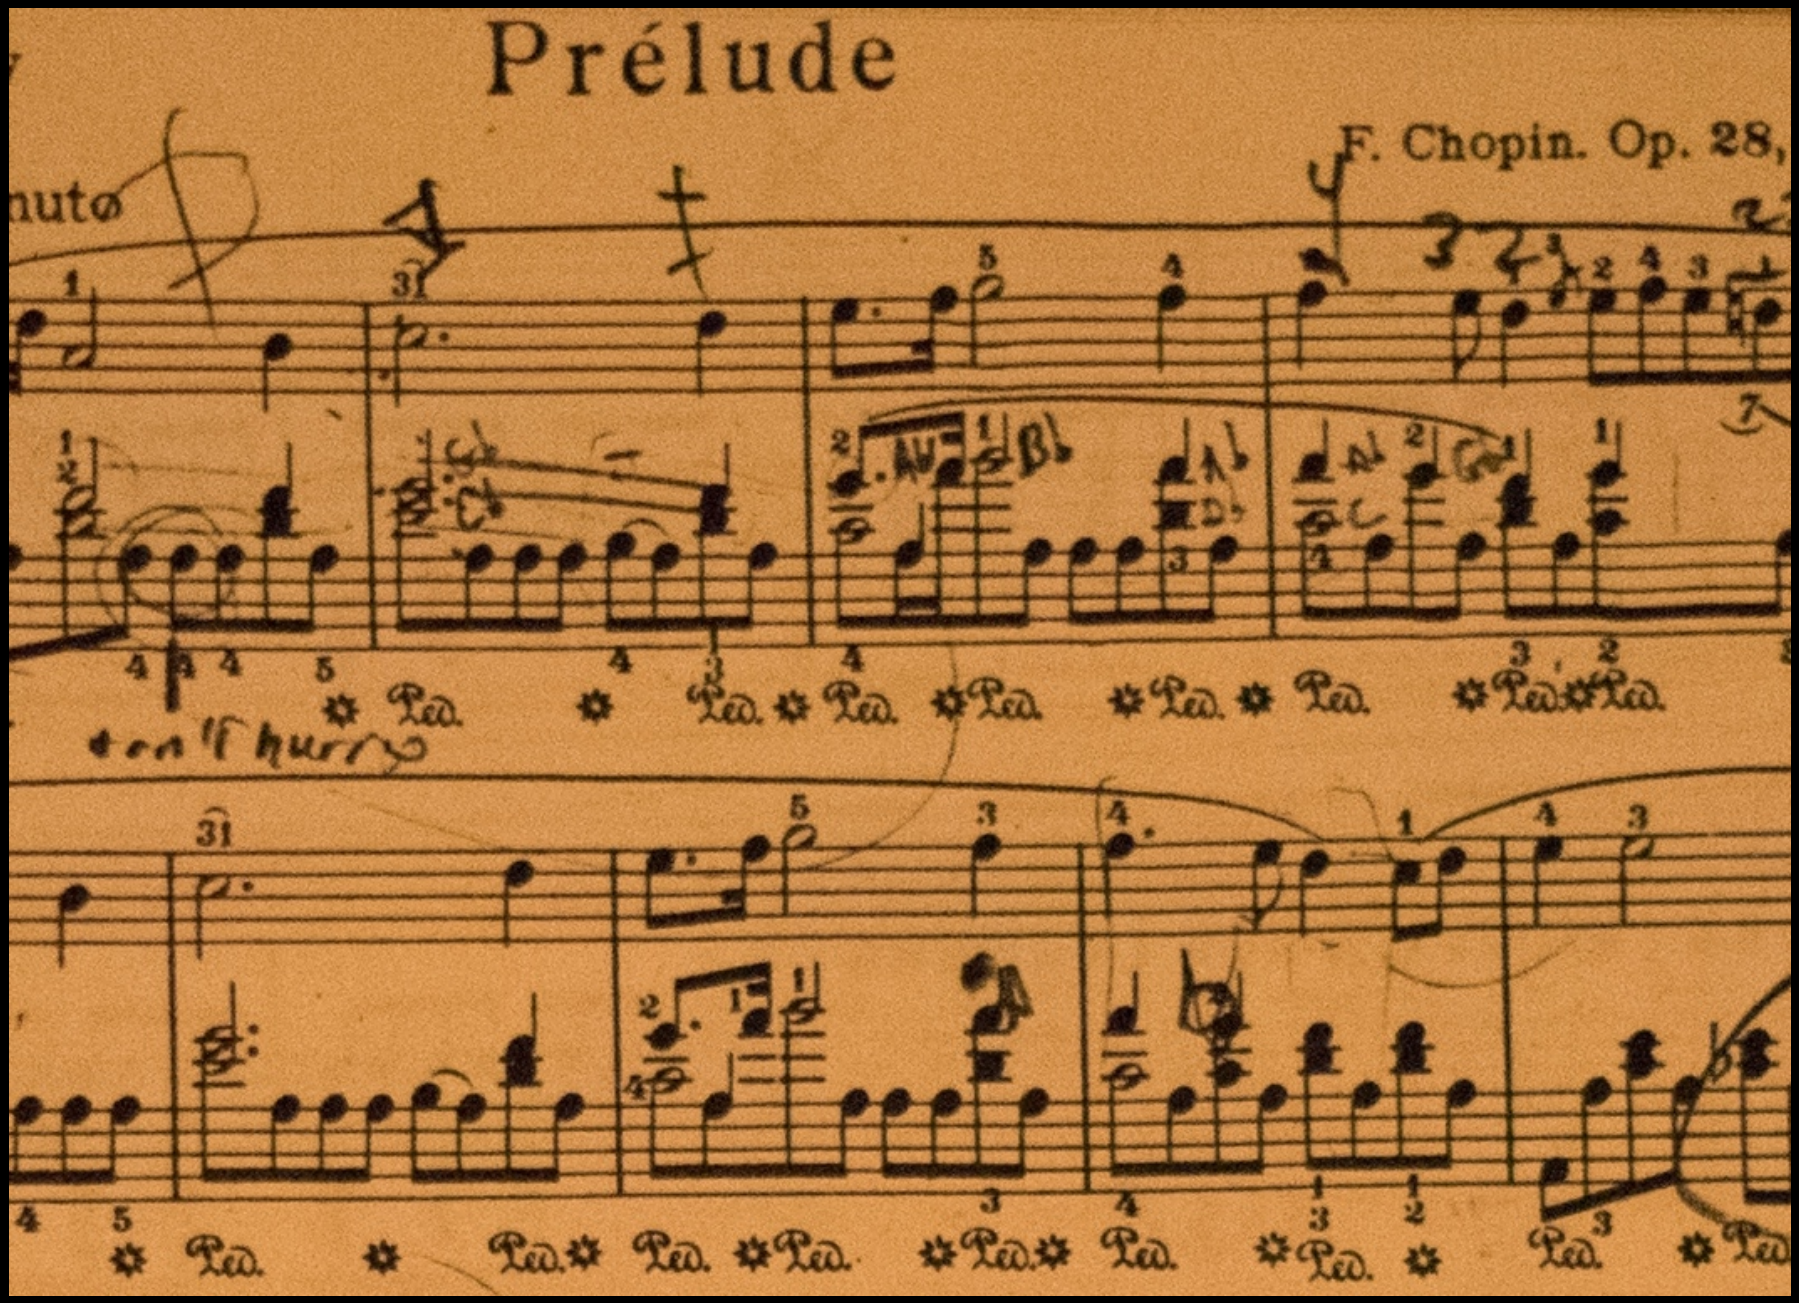

Sony a7II, Leica 35mm prime lens, f/2, ISO 6400,1/8 second

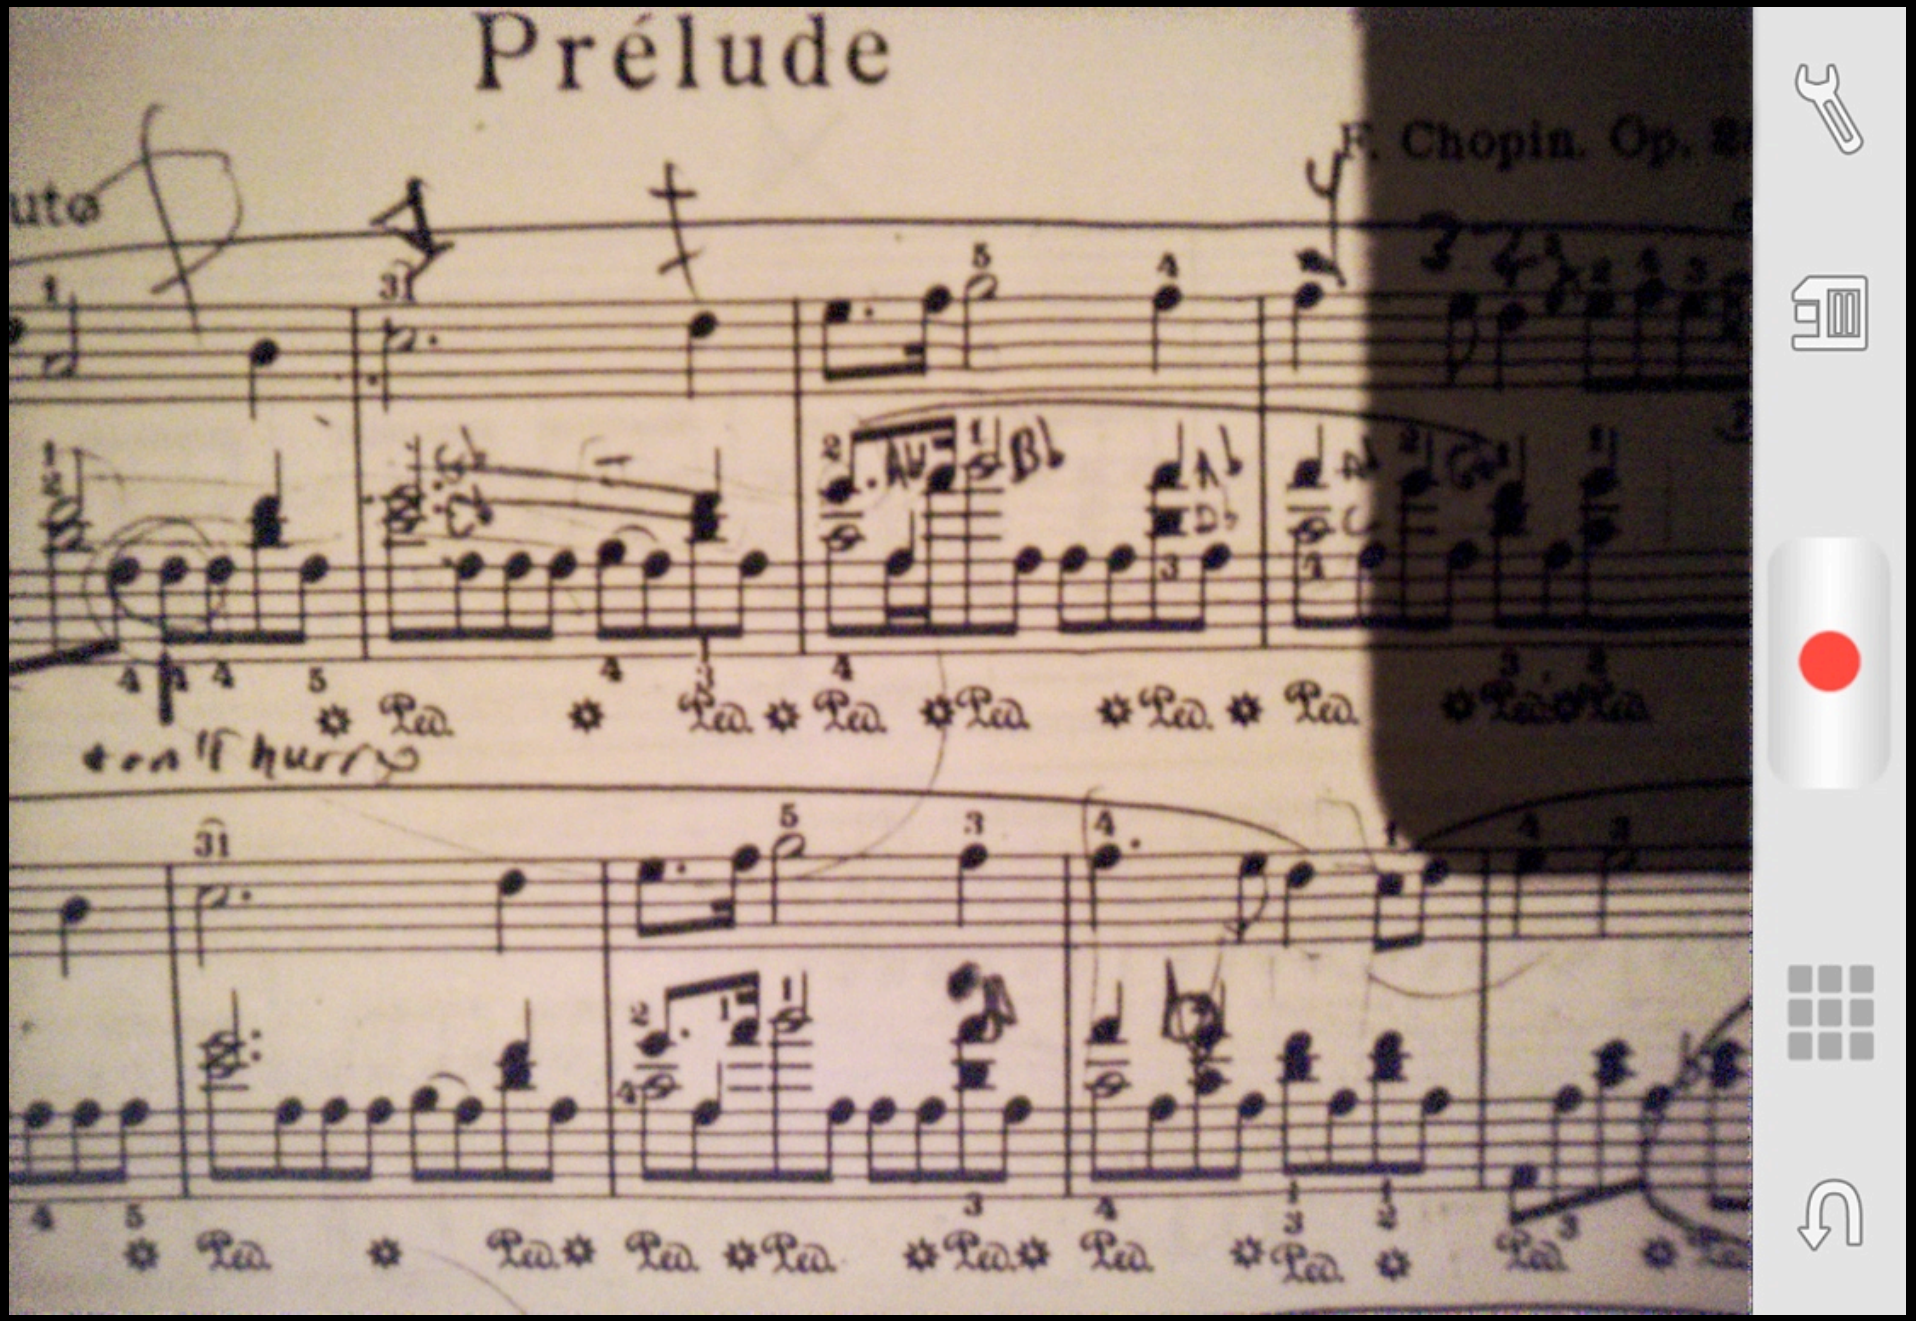

Example computational photography app: HDR+ mode on Nexus phones

Google Research Blog

The latest news from Research at Google

HDR+: Low Light and High Dynamic Range photography in the Google Camera App

Posted: Monday, October 27, 2014

G+1 1.1k

Posted by Marc Levoy, Google[x] Software Engineering Manager and Professor Emeritus, Stanford University

As anybody who has tried to use a smartphone to photograph a dimly lit scene knows, the resulting pictures are often blurry or full of random variations in brightness from pixel to pixel, known as image noise. Equally frustrating are smartphone photographs of scenes where there is a large range of brightness levels, such as a family photo backlit by a bright sky. In high dynamic range (HDR) situations like this, photographs will either come out with an overexposed sky (turning it white) or an underexposed family (turning them into silhouettes).

HDR+ is a feature in the Google Camera app for Nexus 5 and Nexus 6 that uses computational photography to help you take better pictures in these common situations. When you press the shutter button, HDR+ actually captures a rapid burst of pictures, then quickly combines them into one. This improves results in both low-light and high dynamic range situations. Below we delve into each case and describe how HDR+ works to produce a better

Typical approach to HDR

- exposure bracketing
  - capture images with varying exposure
  - combine highlights from short exposure with shadows from long exposure

- hard to robustly handle camera shake or object motion
  - noise level differs between exposures
  - saturated areas cannot be aligned at all

HDR in the Google camera app
[to be published in Proc. SIGGRAPH Asia 2016]

- capture a burst of under-exposed images
  - same exposure on all images in burst
  - avoids blowing out highlights

Stanford Frankencamera architecture and FCam API [Adams SIGGRAPH 2010]

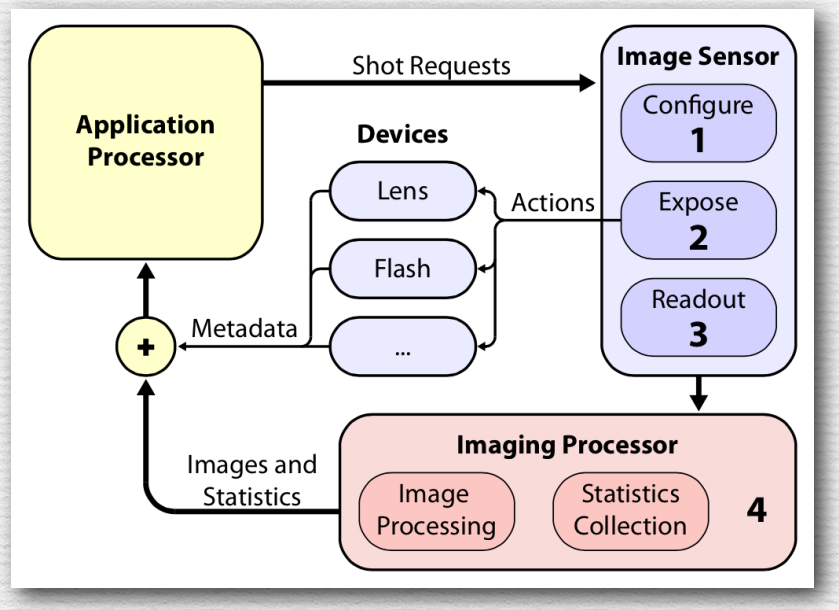


```
Sensor sensor;
Flash flash;
vector<Shot> burst(2);

burst[0].exposure = 1/200.;
burst[1].exposure = 1/30.;

Flash::FireAction fire(&flash);
fire.time = burst[0].exposure/2;
burst[0].actions.insert(fire);

sensor.stream(burst);

while (1) {
  Frame flashFrame =
    sensor.getFrame();
  Frame noflashFrame =
    sensor.getFrame();
}
```

Android Camera HAL 3 architecture and Camera2 API (Eddy Talvala and others)

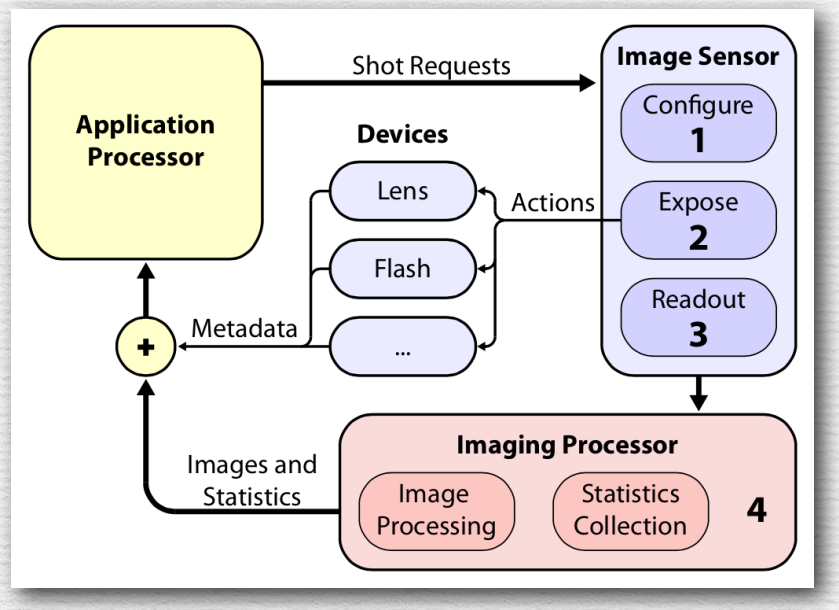


```
Sensor sensor;
Flash flash;
vector<Shot> burst(2);

burst[0].exposure = 1/200.;
burst[1].exposure = 1/30.;

Flash::FireAction fire(&flash);
fire.time = burst[0].exposure/2;
burst[0].actions.insert(fire);

sensor.stream(burst);

while (1) {
  Frame flashFrame =
    sensor.getFrame();
  Frame noflashFrame =
    sensor.getFrame();
}
```

Android Camera HAL 3 architecture and Camera2 API (Eddy Talvala and others)

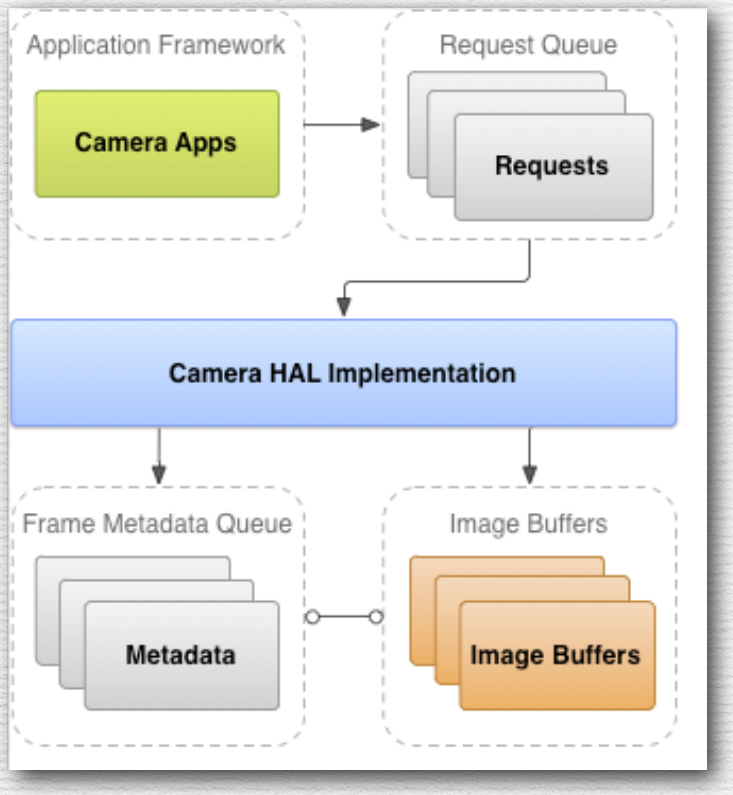


```
Sensor sensor;
Flash flash;
vector<Shot> burst(2);

burst[0].exposure = 1/200.;
burst[1].exposure = 1/30.;

Flash::FireAction fire(&flash);
fire.time = burst[0].exposure/2;
burst[0].actions.insert(fire);

sensor.stream(burst);

while (1) {
  Frame flashFrame =
    sensor.getFrame();
  Frame noflashFrame =
    sensor.getFrame();
}
```

Android Camera HAL 3 architecture and Camera2 API (Eddy Talvala and others)

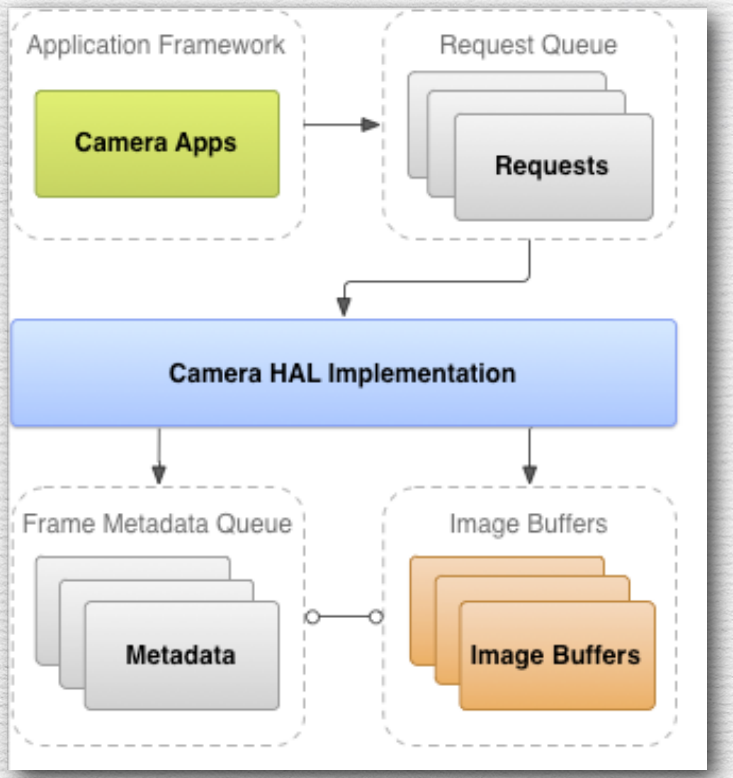

- allows control over the camera
- not over the heterogeneous processors on a mobile platform

open problem!

```
        // This is how to tell the camera to trigger.
        mPreviewRequestBuilder.set(CaptureRequest.CONTROL_AE_PRECAPTURE_TRIGGER,
                CaptureRequest.CONTROL_AE_PRECAPTURE_TRIGGER_START);
        // Tell #mCaptureCallback to wait for the precapture sequence to be set.
        mState = STATE_WAITING_PRECAPTURE;
        mCaptureSession.capture(mPreviewRequestBuilder.build(), mCaptureCallback,
                mBackgroundHandler);
    } catch (CameraAccessException e) {
        e.printStackTrace();
    }
}

/**
 * Capture a still picture. This method should be called when we get a response in
 * {@link #mCaptureCallback} from both {@link #lockFocus()}.
 */
private void captureStillPicture() {
    try {
        final Activity activity = getActivity();
        if (null == activity || null == mCameraDevice) {
            return;
        }
        // This is the CaptureRequest.Builder that we use to take a picture.
        final CaptureRequest.Builder captureBuilder =
                mCameraDevice.createCaptureRequest(CameraDevice.TEMPLATE_STILL_CAPTURE);
        captureBuilder.addTarget(mImageReader.getSurface());

        // Use the same AE and AF modes as the preview.
        captureBuilder.set(CaptureRequest.CONTROL_AF_MODE,
                CaptureRequest.CONTROL_AF_MODE_CONTINUOUS_PICTURE);
        captureBuilder.set(CaptureRequest.CONTROL_AE_MODE,
                CaptureRequest.CONTROL_AE_MODE_ON_AUTO_FLASH);

        // Orientation
        int rotation = activity.getWindowManager().getDefaultDisplay().getRotation();
        captureBuilder.set(CaptureRequest.JPEG_ORIENTATION, ORIENTATIONS.get(rotation));

        CameraCaptureSession.CaptureCallback CaptureCallback
                = new CameraCaptureSession.CaptureCallback() {

            @Override
            public void onCaptureCompleted(CameraCaptureSession session, CaptureRequest request,
                                           TotalCaptureResult result) {
                Toast.makeText(getActivity(), "Saved: " + mFile, Toast.LENGTH_SHORT).show();
                unlockFocus();
            }
        };

        mCaptureSession.stopRepeating();
        mCaptureSession.capture(captureBuilder.build(), CaptureCallback, null);
    } catch (CameraAccessException e) {
        e.printStackTrace();
    }
}
```

HDR in the Google camera app
[to be published in Proc. SIGGRAPH Asia 2016]

- capture a burst of under-exposed images
  - same exposure on all images in burst
  - avoids blowing out highlights

- align and merge
  - all images look similar, so alignment is more robust
  - SNR ∝ sqrt(N) where N = number of frames in burst
  - reduces noise in shadows

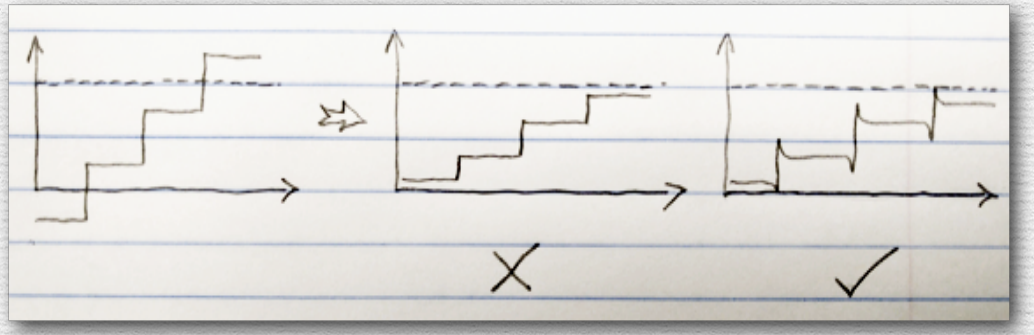

- tonemap
  - boost shadows
  - squeeze 14-bit merged image into 8-bit for display
  - preserve local contrast at the expense of global contrast

HDR in the Google camera app

[to be published in Proc. SIGGRAPH Asia 2016]

- capture a burst of under-exposed images
  - same exposure on all images in burst
  - avoids blowing out highlights
- align and merge
  - all images look similar, so alignment is more robust
  - reduces noise in shadows
- tonemap
  - boost shadows
  - enhance local contrast to avoid washed-out look

HDR in the Google camera app

[to be published in Proc. SIGGRAPH Asia 2016]

the Mach band illusion:
each wedge should appear
brighter on its right side

- align and merge
  - all images look similar, so alignment is more robust
  - SNR ∝ sqrt(N) where N = number of frames in burst
  - reduces noise in shadows

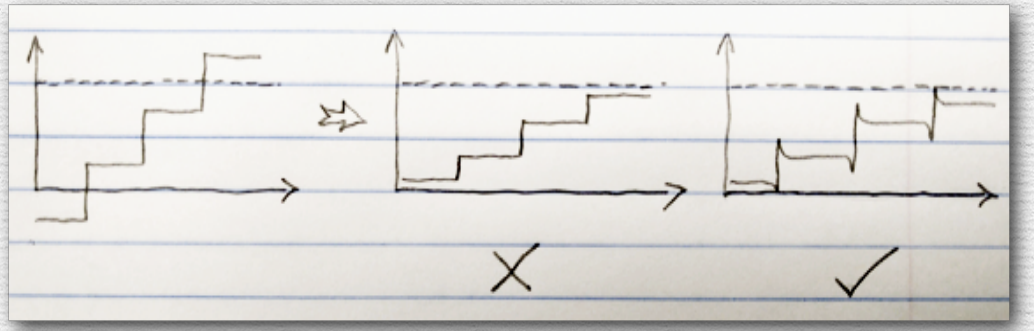

- tonemap
  - boost shadows
  - squeeze 14-bit merged image into 8-bit for display
  - preserve local contrast at the expense of global contrast

HDR in the Google camera app

[to be published in Proc. SIGGRAPH Asia 2016]

the Mach band illusion:
each wedge should appear
brighter on its right side

- align and merge
  - all images look similar, so alignment is more robust
  - SNR ∝ sqrt(N) where N = number of frames in burst
  - reduces noise in shadows

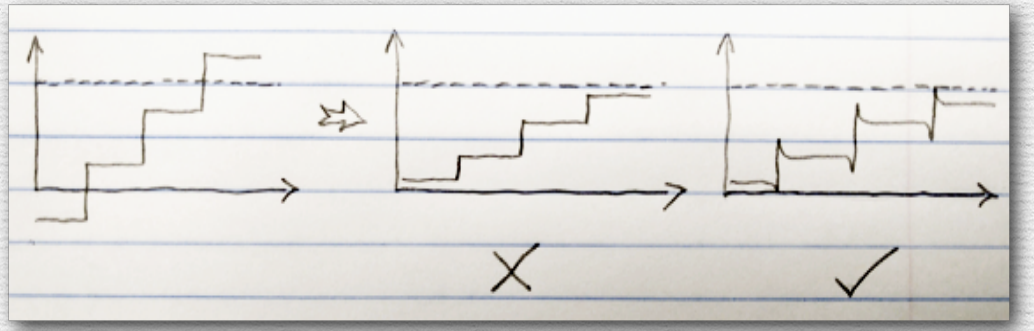

- tonemap
  - boost shadows
  - squeeze 14-bit merged image into 8-bit for display
  - preserve local contrast at the expense of global contrast

La Grande Jatte, Georges Seurat, 1884

HDR in the Google camera app

[to be published in Proc. SIGGRAPH Asia 2016]

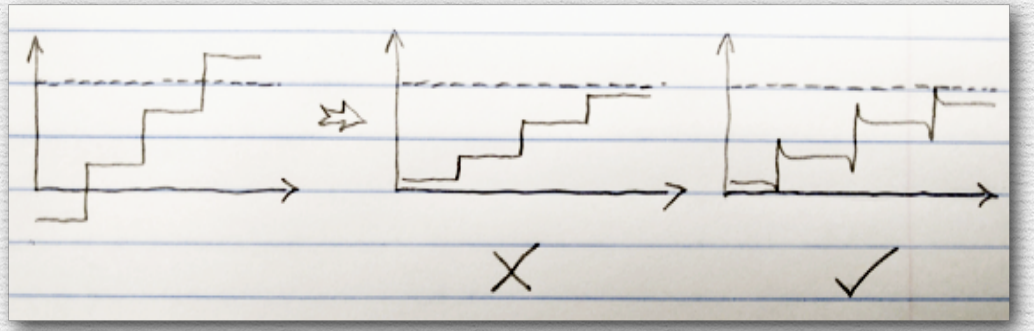

- tonemap
  - boost shadows
  - squeeze 14-bit merged image into 8-bit for display
  - preserve local contrast at the expense of global contrast


single frame


HDR+


single frame


HDR+


single frame


HDR+

23 mins after sunset

Nexus 6

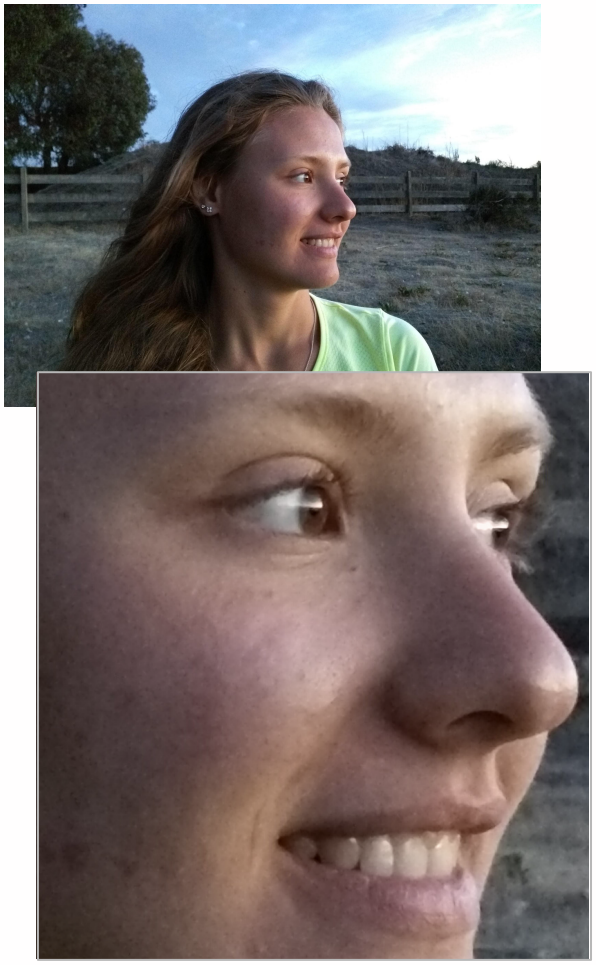

iPhone 6S+

Nexus 6P

37 mins after sunset

Nexus 6

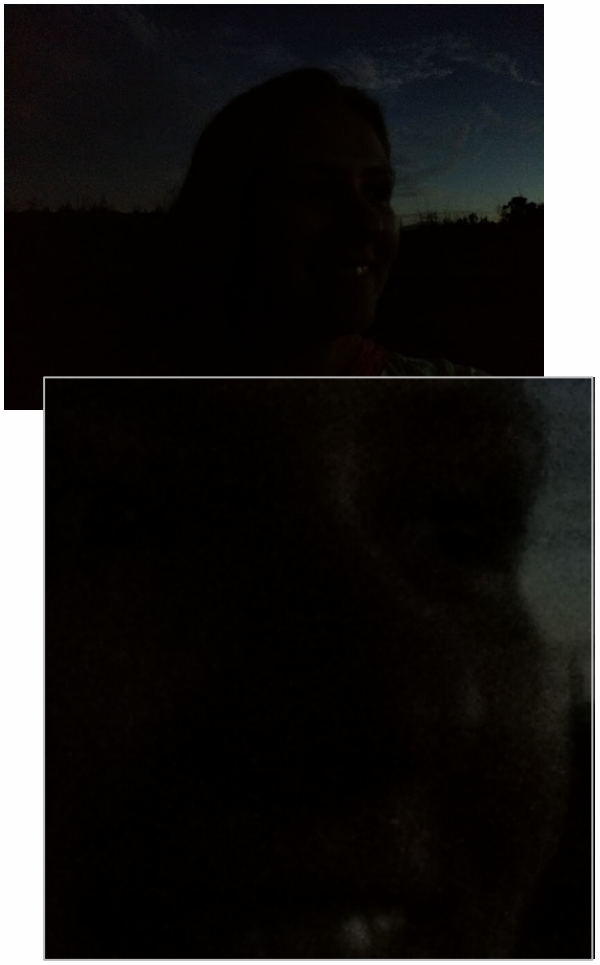

iPhone 6S+

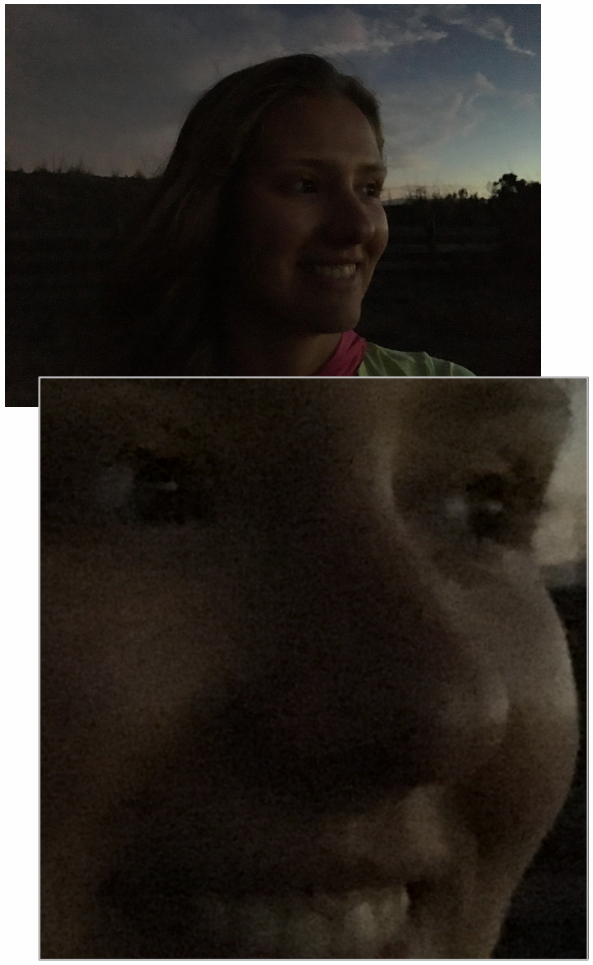

Nexus 6P

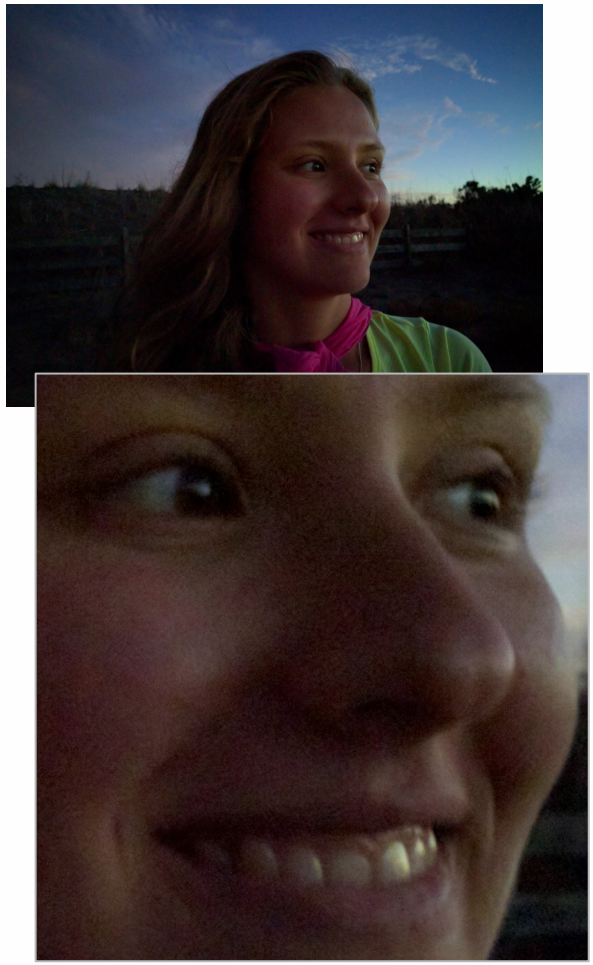


single frame


HDR+

WHAT IF THE SCENE
WERE DARKER?

(apologies to xkcd)

Illuminance on a surface

- power accumulating on a surface per unit area, considering light arriving from all directions

$$E = \frac{P}{A} \qquad \left(\frac{\text{lumens}}{\text{m}^2}\right)$$

(Reinhard)

Illuminance on a surface

- power accumulating on a surface per unit area, considering light arriving from all directions

$$E = \frac{P}{A} \qquad \left(\frac{\text{lumens}}{\text{m}^2} \right)$$

- related units
  - 1 lux = 1 lumen / m^2
  - British unit is footcandle
    - 1 candela held 1 foot from surface
    - 1 footcandle = 10.764 lux

Lux levels

300 lux = office lighting at Google

150 lux = desk lighting at home

50 lux = average restaurant

20 lux = romantic restaurant

10 lux = finding socks that match

Lux levels

300 lux = office lighting at Google

150 lux = desk lighting at home

50 lux = average restaurant

20 lux = romantic restaurant

10 lux = finding socks that match

3 lux = outdoor street lighting

1 lux = limit of reading

1/2 lux = full moon

1/5 lux = can't find my keys

1/10 lux = wouldn't take a step without a flashlight

1/4 lux, iPhone 6S+

HDR+

SeeInTheDark, ~50 frames, handheld, real-time

Operation of SeeInTheDark

- capture 10-bit Bayer RAW frame
  - Nexus 6P capable of 12Mpix at 30fps, but I use only 8fps
- minimal pre-processing
  - correct for fixed pattern noise and hot pixels
  - downsize 4× in X and Y (instead of demosaicing) to 1Mpix

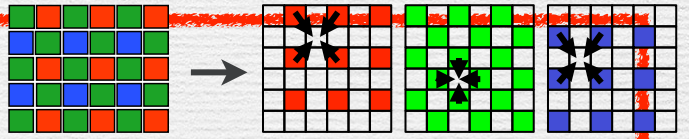

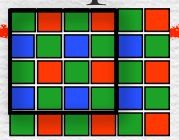

- track feature(s) from previous frame
  - 2D convolutions, then fit homography, affine, similarity, or simple translation, with help from accelerometer and gyro
- combine in GPU using 4 × 16-bit FBOs (for R,G,B,G)
  - warp new frame to accumulator and blend in using IIR filter
  - apply black level correction, lens shading, white balancing, color correction, tone mapping, gamma curve, sharpening

Fixed pattern noise

- differences in readout among columns
- rises with ISO but independent of exposure time

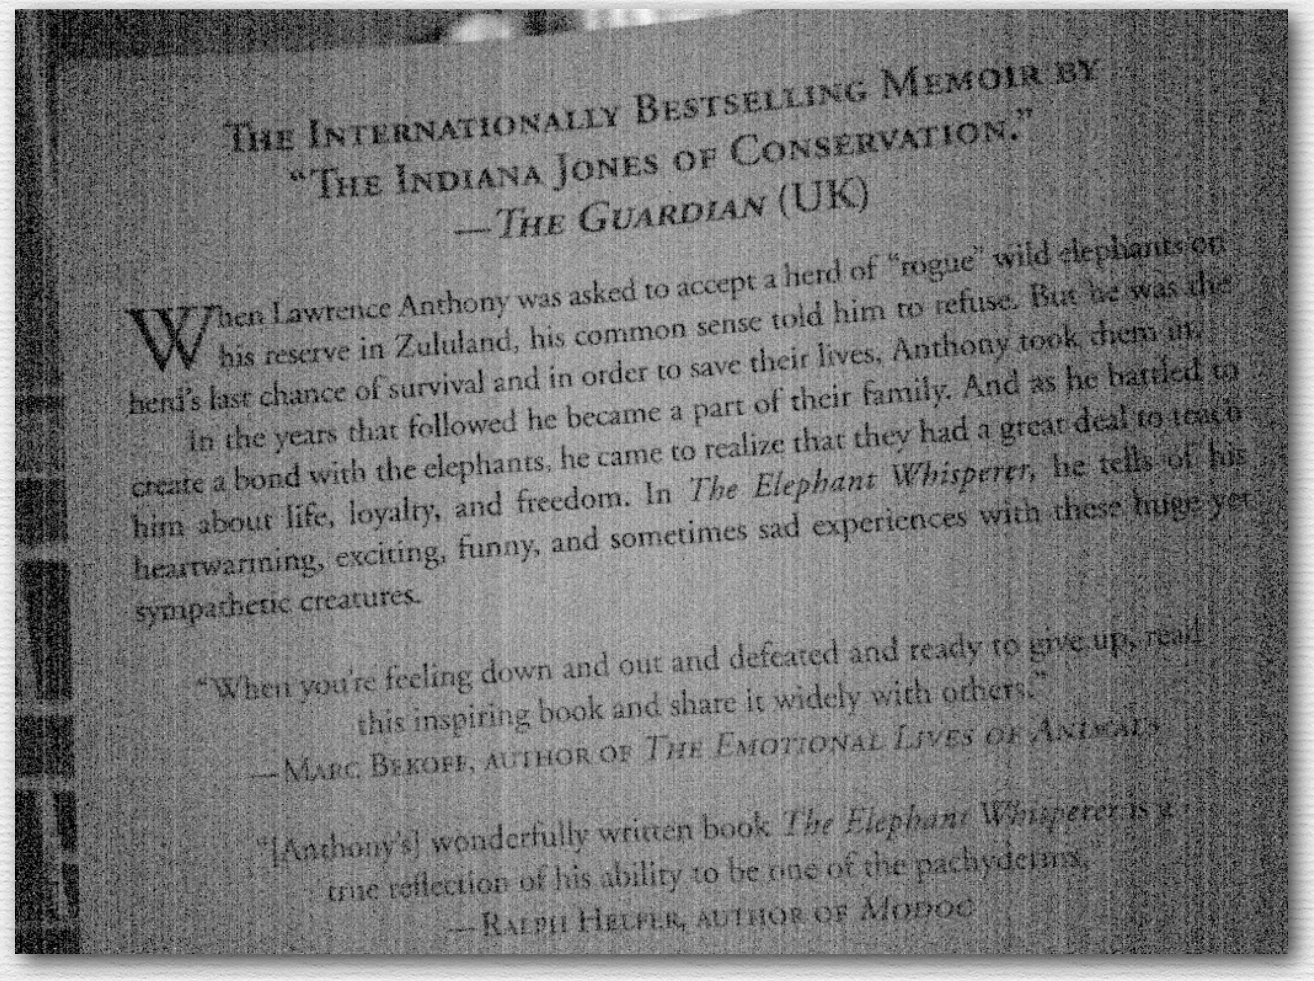

Fixed pattern noise

- differences in readout among columns
- rises with ISO but independent of exposure time
- measure using short exposure, but one shot is too noisy

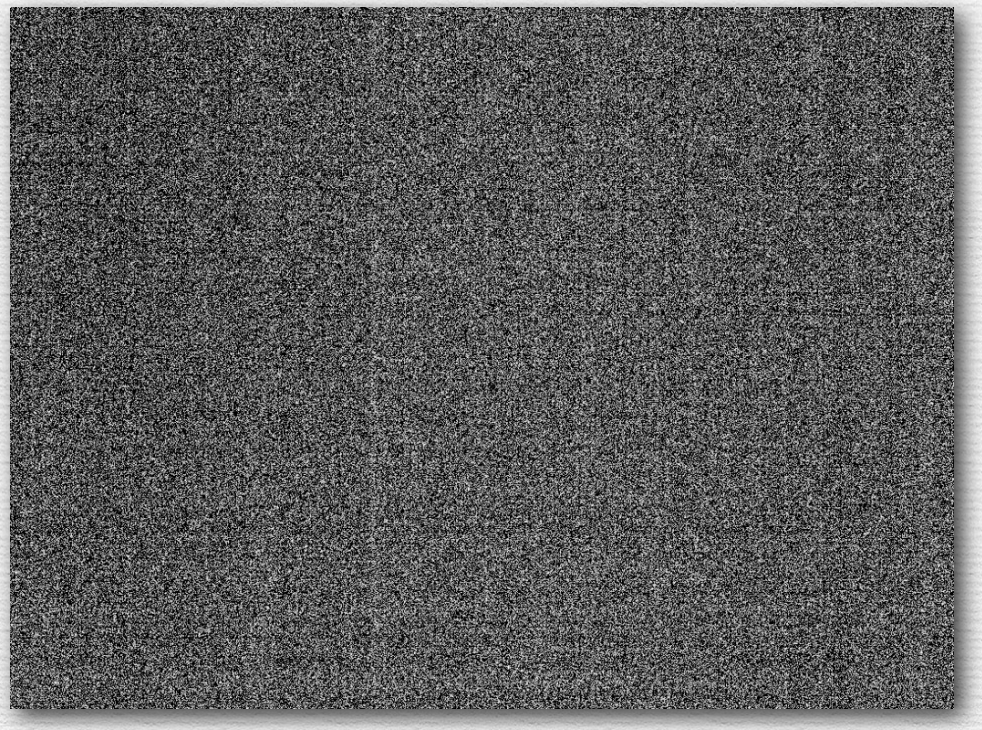

Fixed pattern noise

- differences in readout among columns
- rises with ISO but independent of exposure time
- measure using short exposure, but one shot is too noisy
- accumulate many short exposures to build a corrector

THE INTERNATIONALLY BESTSELLING MEMOIR BY
"THE INDIANA JONES OF CONSERVATION."
—*THE GUARDIAN* (UK)

When Lawrence Anthony was asked to accept a herd of "rogue" wild elephants on his reserve in Zululand, his common sense told him to refuse. But he was the herd's last chance of survival and in order to save their lives, Anthony took them all in the years that followed he became a part of their family. And as he battled to create a bond with the elephants, he came to realize that they had a great deal to teach him about life, loyalty, and freedom. In *The Elephant Whisperer*, he tells of his heartwarming, exciting, funny, and sometimes sad experiences with these huge yet sympathetic creatures.

"When you're feeling down and out and defeated and ready to give up, read this inspiring book and share it widely with others."
—MARC BEKOFF, AUTHOR OF *THE EMOTIONAL LIVES OF ANIMALS*

"[Anthony's] wonderfully written book *The Elephant Whisperer* is a true reflection of his ability to be one of the pachyderms."
—RALPH HELFER, AUTHOR OF *MODOC*

THE INTERNATIONALLY BESTSELLING MEMOIR BY
"THE INDIANA JONES OF CONSERVATION."
—*THE GUARDIAN* (UK)

When Lawrence Anthony was asked to accept a herd of "rogue" wild elephants on his reserve in Zululand, his common sense told him to refuse. But he was the herd's last chance of survival and in order to save their lives, Anthony took them in. In the years that followed he became a part of their family. And as he battled to create a bond with the elephants, he came to realize that they had a great deal to teach him about life, loyalty, and freedom. In *The Elephant Whisperer,* he tells of his heartwarming, exciting, funny, and sometimes sad experiences with these huge yet sympathetic creatures.

"When you're feeling down and out and defeated and ready to give up, read this inspiring book and share it widely with others."
—MARC BEKOFF, AUTHOR OF *THE EMOTIONAL LIVES OF ANIMALS*

"[Anthony's] wonderfully written book *The Elephant Whisperer* is a true reflection of his ability to be one of the pachyderms."
—RALPH HELFER, AUTHOR OF *MODOC*

Hot pixels

- isolated pixels with excessive current leakage
- rises with ISO and exposure time

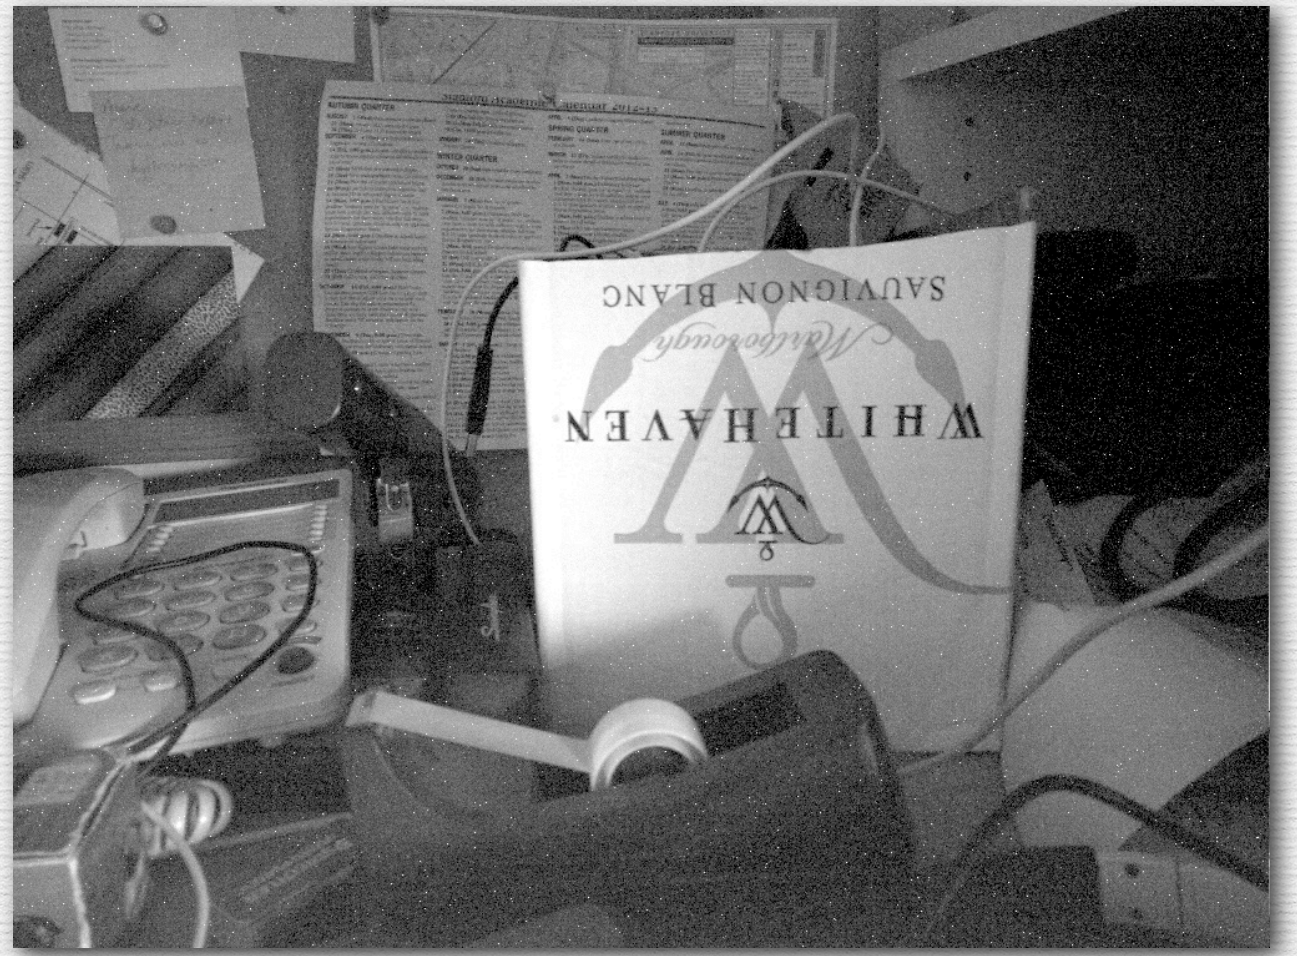

Removing hot pixels

- just look for isolated bright pixels, right? No, too noisy!

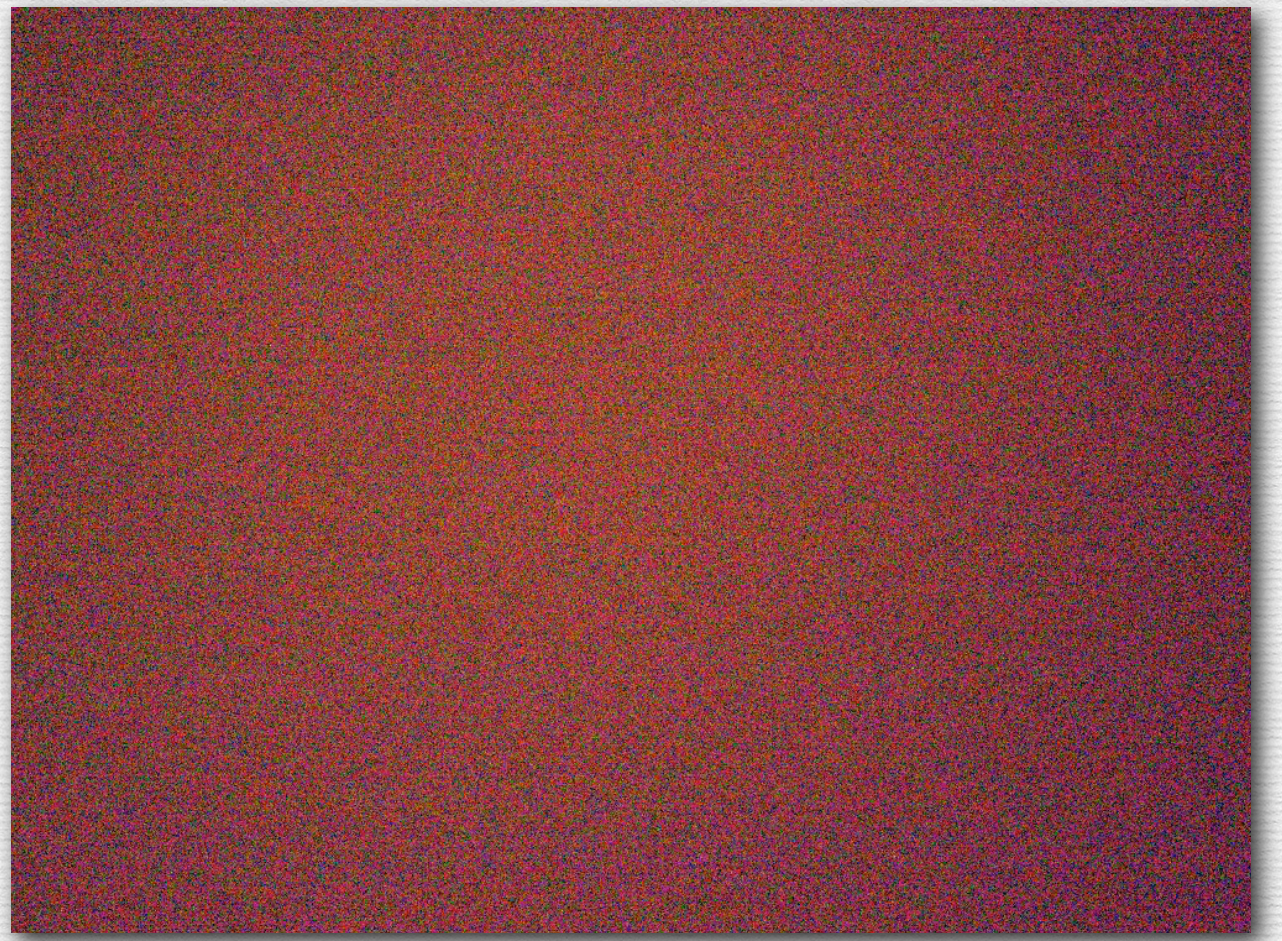

Removing hot pixels

- just look for isolated bright pixels, right? No, too noisy!
- accumulate to reduce noise, then look for them

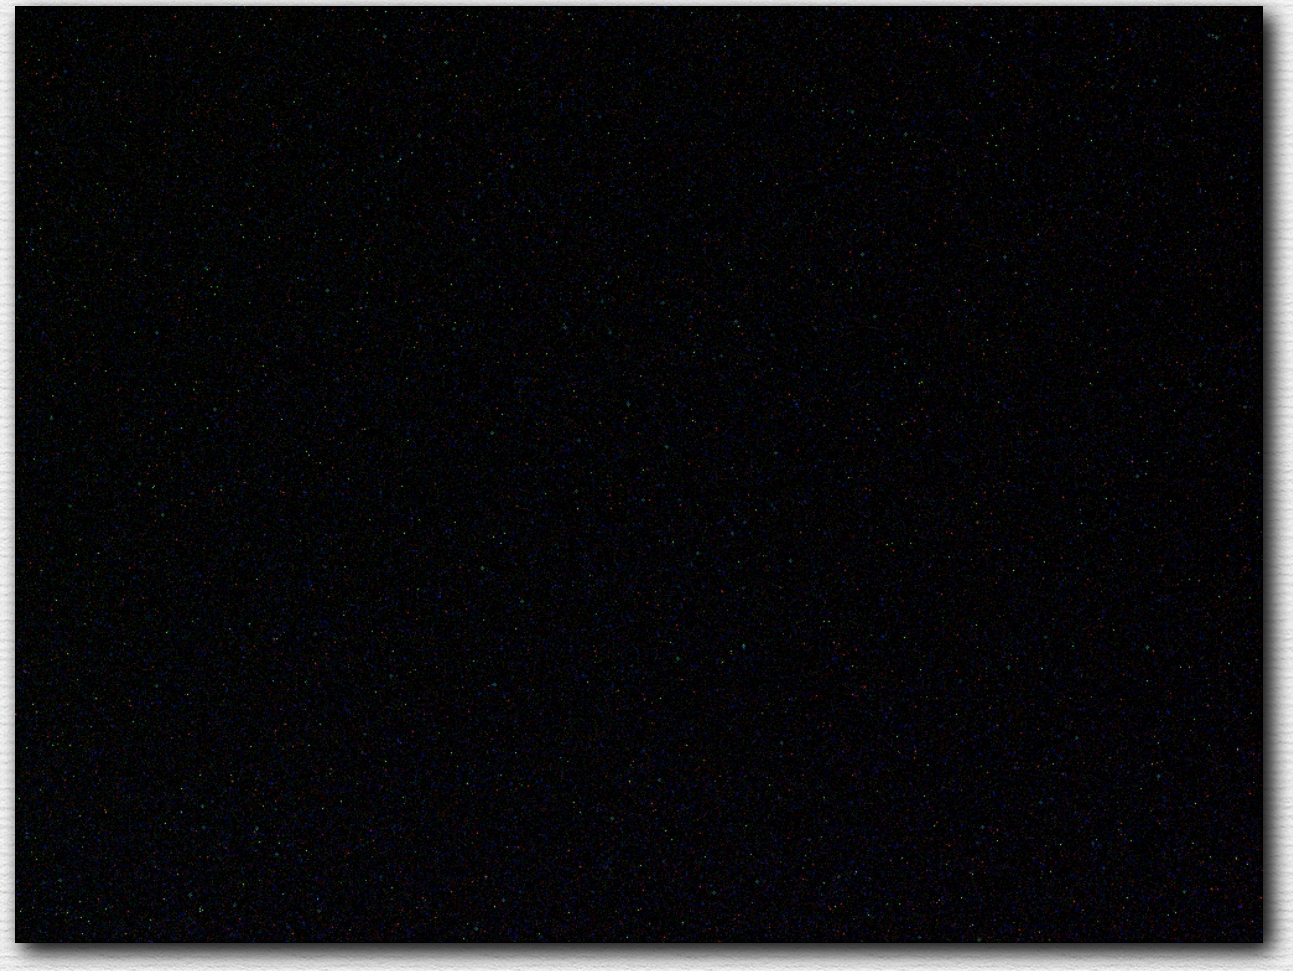

Removing hot pixels

- just look for isolated bright pixels, right? No, too noisy!
- accumulate to reduce noise, then look for them

Q. Need long exposure, but cell phones have no shutters.
If you see the scene, you'll be fooled by sharp features.
What to do?

A. Accumulate during auto-focus sweep!

Removing hot pixels

- just look for isolated bright pixels, right? No, too noisy!
- accumulate to reduce noise, then look

  Q. Need long exposure, but cell phones have no shutters.
  If you see the scene, you'll be fooled by sharp features.
  What to do?

  A. Accumulate during auto-focus sweep!

- replace hot pixels with average of its neighbors

before hot pixel correction

after hot pixel correction

Operation of SeeInTheDark

- capture 10-bit Bayer RAW frame
  - Nexus 6P capable of 12Mpix at 30fps, but I use only 8fps
- minimal pre-processing
  - correct for fixed pattern noise and hot pixels
  - downsize 4× in X and Y (instead of demosaicing) to 1Mpix
- track feature(s) from previous frame
  - 2D convolutions, then fit homography, affine, similarity, or simple translation, with help from accelerometer and gyro
- combine in GPU using 4 × 16-bit FBOs (for R,G,B,G)
  - warp new frame to accumulator and blend in using IIR filter
  - apply black level correction, lens shading, white balancing, color correction, tone mapping, gamma curve, sharpening

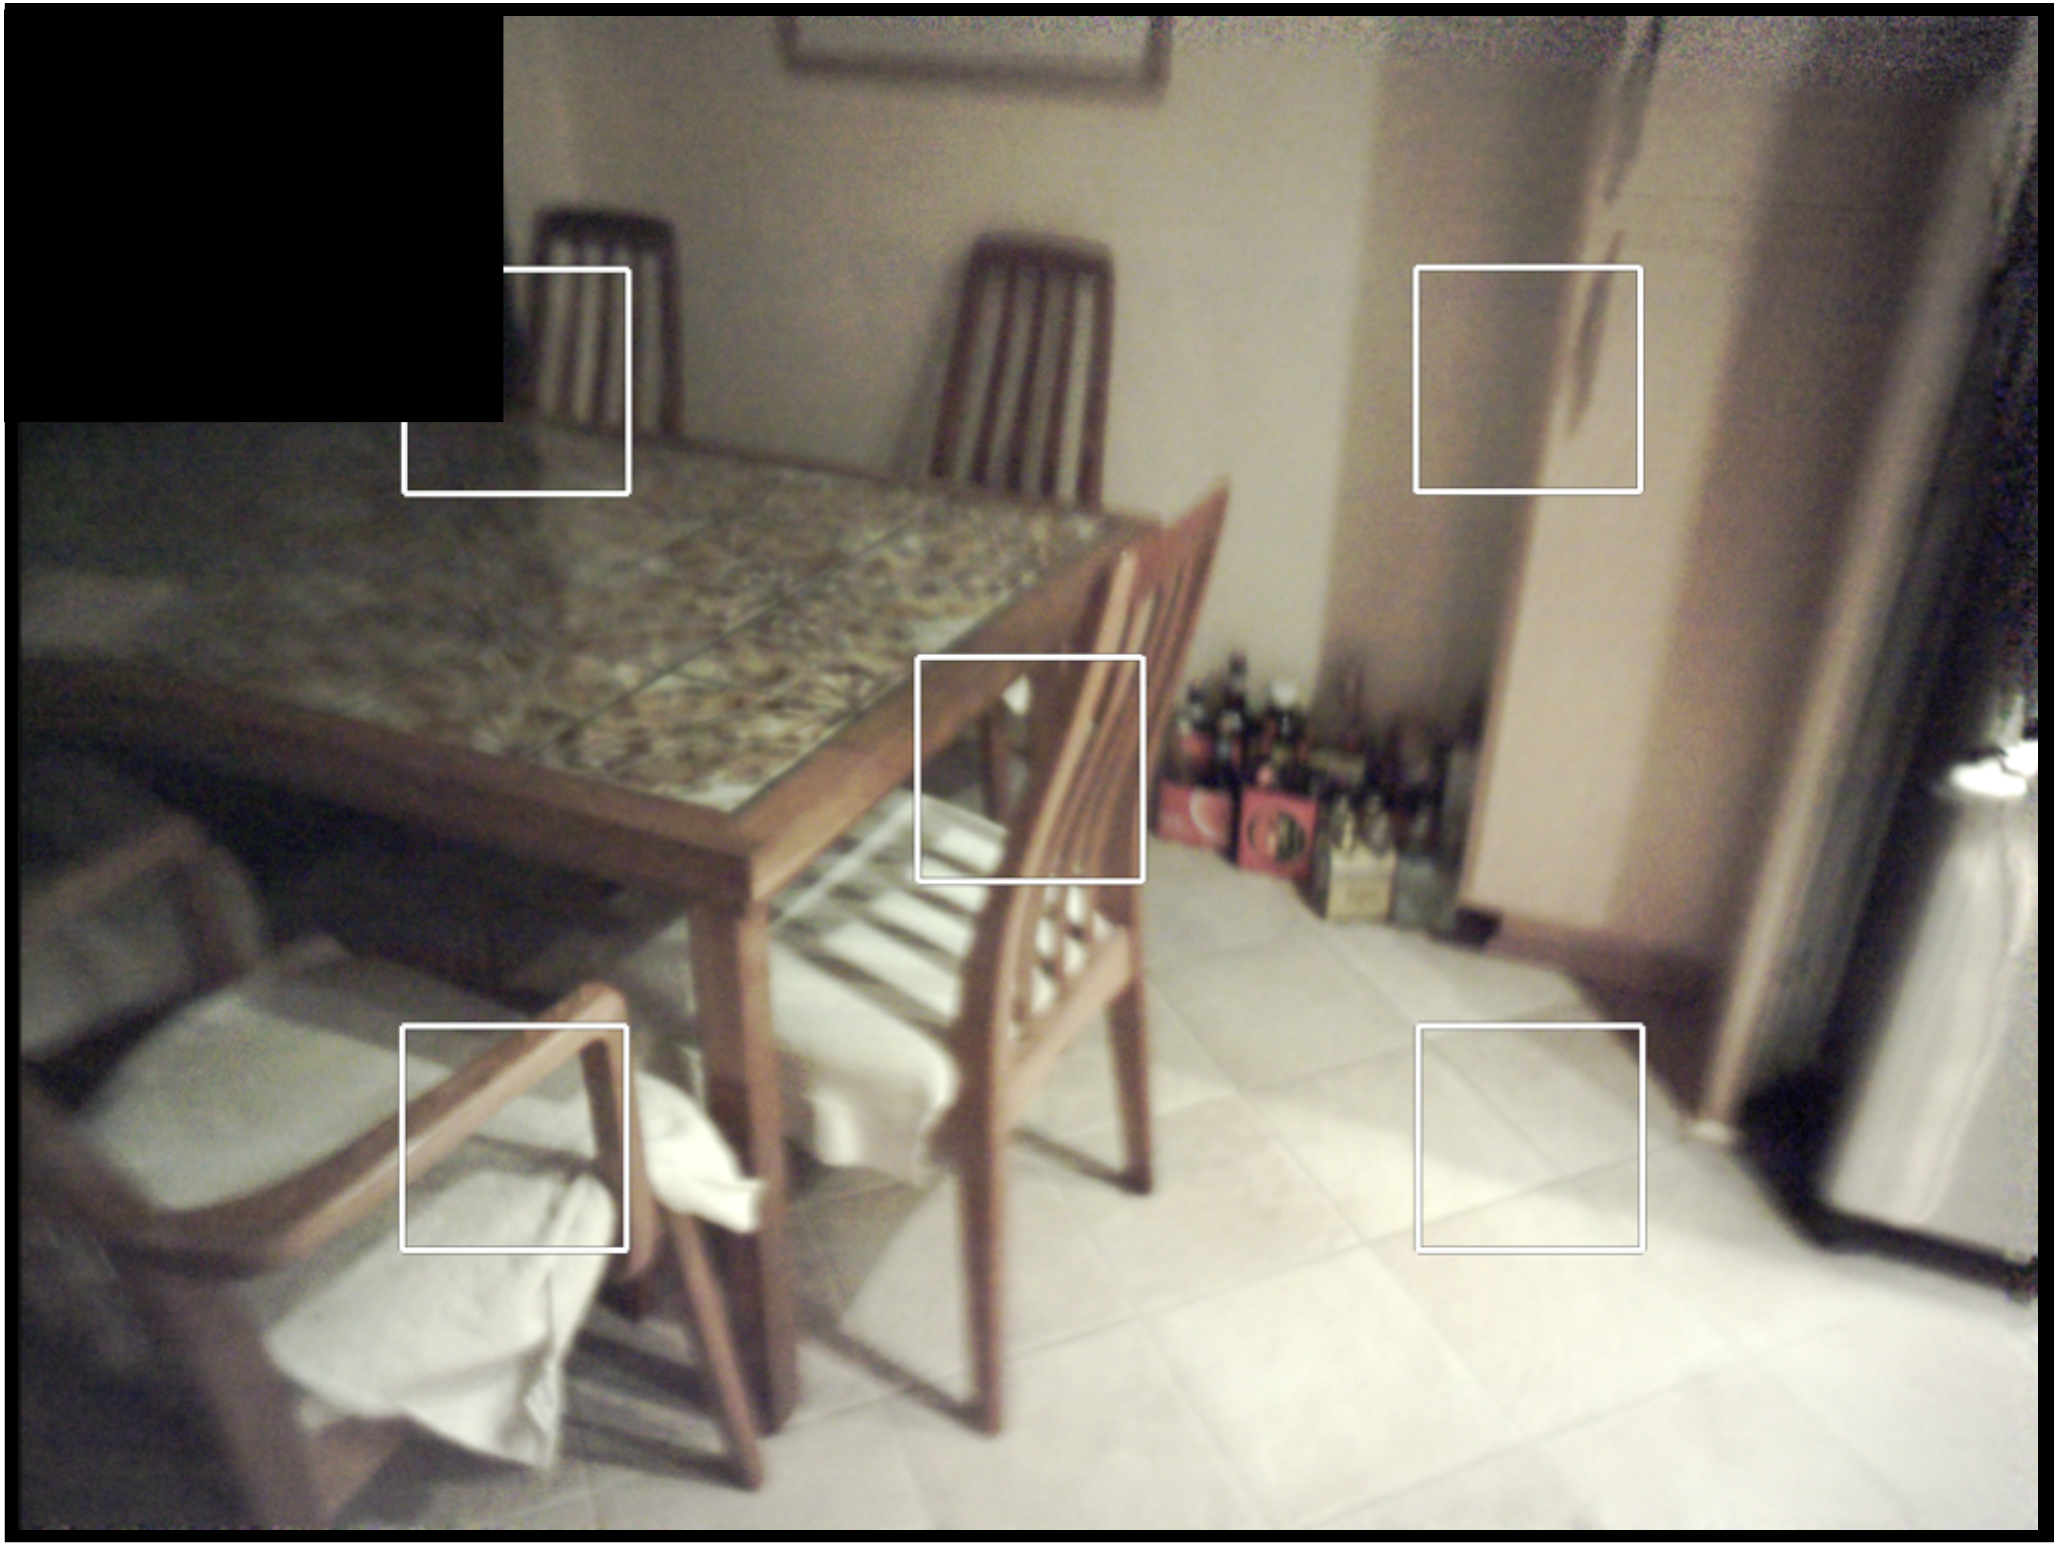

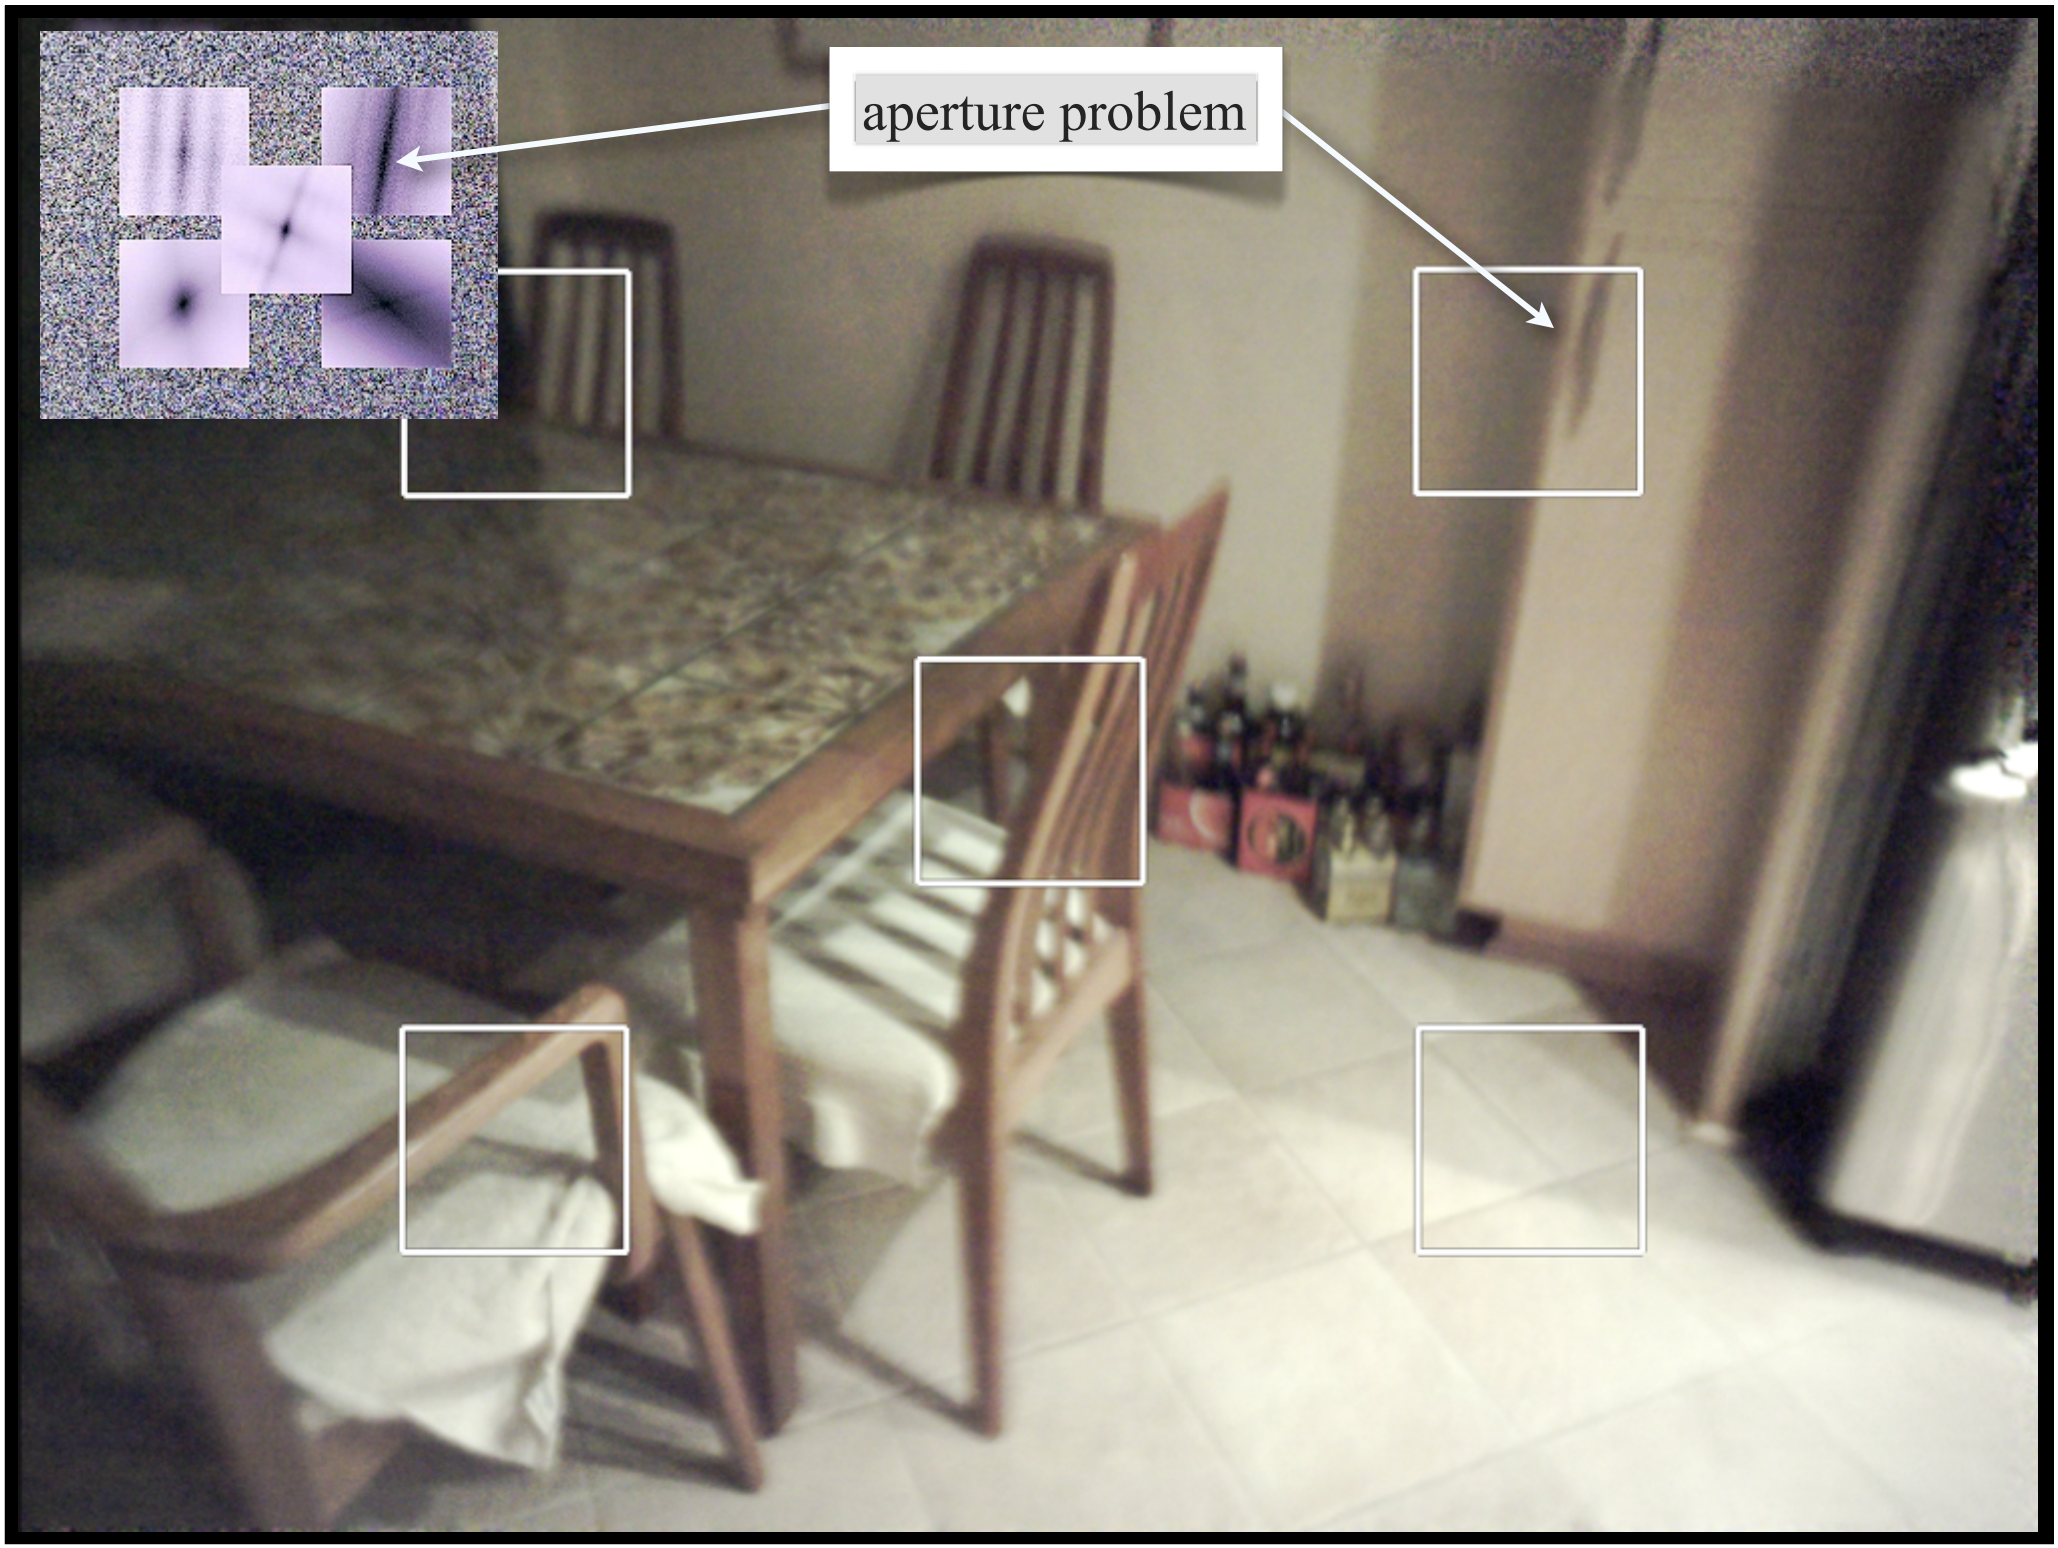
aperture problem


2213
PACIFIC INTERNET


vignetting causes
loss of signal
2213
PACIFIC INTERNET www.pacific.net
DEATH RIDE
YELLOW CAB CO.


2213
PACIFIC INTERNET

[illegible] Creek Open Space Preserve

California Buckeye in Winter

Sunday • January 17
Rancho [illegible] • 4.5 mi
9:30am to 2:30pm

Explore the buckeye trees of Black Mountain with Docents Mike Alexander and Kandis [illegible] on a moderately-paced hike along (and adjacent to) the Bella Vista and Old Ranch Trails. You'll learn about this remarkable species' history, uses, and unique adaptations to California's climate. In winter, the buckeye exhibits a stark beauty like no other, and you'll observe [illegible] variety of inspiring, and [illegible] settings. You'll [illegible] [illegible] back then [illegible]

[illegible] climb Black Mountain, ride [illegible] Creek Trail loop, and [illegible] on rider preferences and [illegible] continue to White Oak [illegible]. Reservations are required and will be accepted on or after January 5.

Bay Breezes and Windy [illegible]

Tuesday • January 19
Windy Hill • 5 mi
10:00am to 3:00pm

Join Docents Ed [illegible] [illegible] for a tour of this Preserve [illegible] how it got its name. You'll [illegible] Gulch and traverse [illegible] weather [illegible] [illegible]

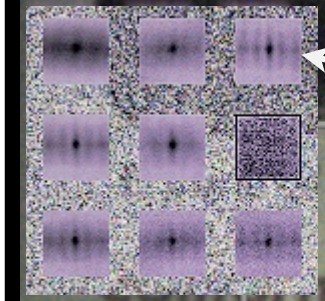

text is a repeating texture

Madera Creek Open S

...nia Buckeye in Winter

Sunday • January 17
Monte Bello • 4.5 mi
9:30am to 2:30pm

Explore the buckeye trees of Black Mountain with docents Mike Alexander and Kandis Scott on a moderately-paced hike along (and adjacent to) the Bella Vista and Old Ranch Trails. You'll learn about this remarkable species' history, uses, and unique adaptations to California's climate.

In winter, the buckeye exhibits a stark beauty like no other, and you'll observe this in a variety of inspiring, and photogenic forms and settings. You'll help record detailed observations of select trees as a way to track their seasonal progression. This is the first

not to go at a faster cycling pace. You'll ...Mountain, ride the Waterwheel ...p, and perhaps (depending on rider preferences and trail availability) continue to White Oak Trail to climb out. **Reservations are required and will be accepted on or after January 3.**

Bay Breezes and Windy Hills

Tuesday • January 19
Windy Hill • 8 mi
10:00am to 3:00pm

Join Docents Ed North and Huey-Shin Yuan for a tour of this Preserve and learn how it got its name. You'll climb Hamms Gulch and traverse the Lost Trail, and if the weather cooperates, you may be able to get a spectacular glimpse of the Bay. You'll return on the Spring Ridge Trail on this strenuous hike that includes approximately 1,200 feet of climbing.

[illegible] Creek Open Space Preserve

California Buckeye in Winter

Sunday • January 17
Monte Bello • 4.5 mi
9:30am to 2:30pm

Explore the buckeye trees of Black Mountain with Docents Mike Alexander and Kandis Scott on a moderately-paced hike along (and adjacent to) the Bella Vista and Old Ranch Trails. You'll learn about this remarkable species' history, uses, and unique adaptations to California's climate. In winter, the buckeye exhibits a stark beauty like no other, and you'll observe them in a variety of inspiring, and [illegible] settings. You'll [illegible]

[illegible] climb Black Mountain, [illegible] Creek Trail loop, and [illegible] on rider preferences and [illegible] continue to White Oak Trail [illegible]. Reservations are required and will be accepted on or after January 5.

Bay Breezes and Windy Hills

Tuesday • January 19
Windy Hill • 8 mi
10:00am to 3:00pm

Join Docents Ed [illegible] Yuan for a tour of this Preserve [illegible] how it got its name. You'll [illegible] Gulch and traverse [illegible] weather [illegible]

How much can the gyro help?

- MEMS gyros are accurate, but they drift
  - ~1 pixel / frame (for 1Mpix image@8fps)
- yaw is worst, because it can't be corrected using the gravity vector from the accelerometer
- the compass is too inaccurate to correct anything

after 1 second

after 15 seconds

Operation of SeeInTheDark

- capture 10-bit Bayer RAW frame
  - Nexus 6P capable of 12Mpix at 30fps, but I use only 8fps
- minimal pre-processing
  - correct for fixed pattern noise and hot pixels
  - downsize 4× in X and Y (instead of demosaicing) to 1Mpix
- track feature(s) from previous frame
  - 2D convolutions, then fit homography, affine, similarity, or simple translation, with help from accelerometer and gyro
- combine in GPU using 4 × 16-bit FBOs (for R,G,B,G)
  - warp new frame to accumulator and blend in using IIR filter
  - apply black level correction, lens shading, white balancing, color correction, tone mapping, gamma curve, sharpening

accumulator after first frame

second frame

second frame warped to match accumulator

accumulator after blending in second frame


1 lux
after blending in many more frames

Two ways to average images

- FIR filter

$$y[n] = b_0\,x[n] + b_1\,x[n\text{-}1] + \ldots + b_N\,x[n\text{-}N]$$

  - Example:

$$\text{accum}[n] = \tfrac{1}{3}\,\text{frame}[n] + \tfrac{1}{3}\,\text{frame}[n\text{-}1] + \tfrac{1}{3}\,\text{frame}[n\text{-}2]$$

O(N) memory

- IIR filter

$$y[n] = k_1\,x[n] + k_2\,y[n\text{-}1]$$

  - Example:

$$\text{accum}[n] = \tfrac{1}{2}\,\text{frame}[n] + \tfrac{1}{2}\,\text{accum}[n\text{-}1]$$
$$= \tfrac{1}{2}\,\text{frame}[n] + \tfrac{1}{4}\,\text{frame}[n\text{-}1] + \tfrac{1}{8}\,\text{frame}[n\text{-}2] + \ldots$$

O(1) memory

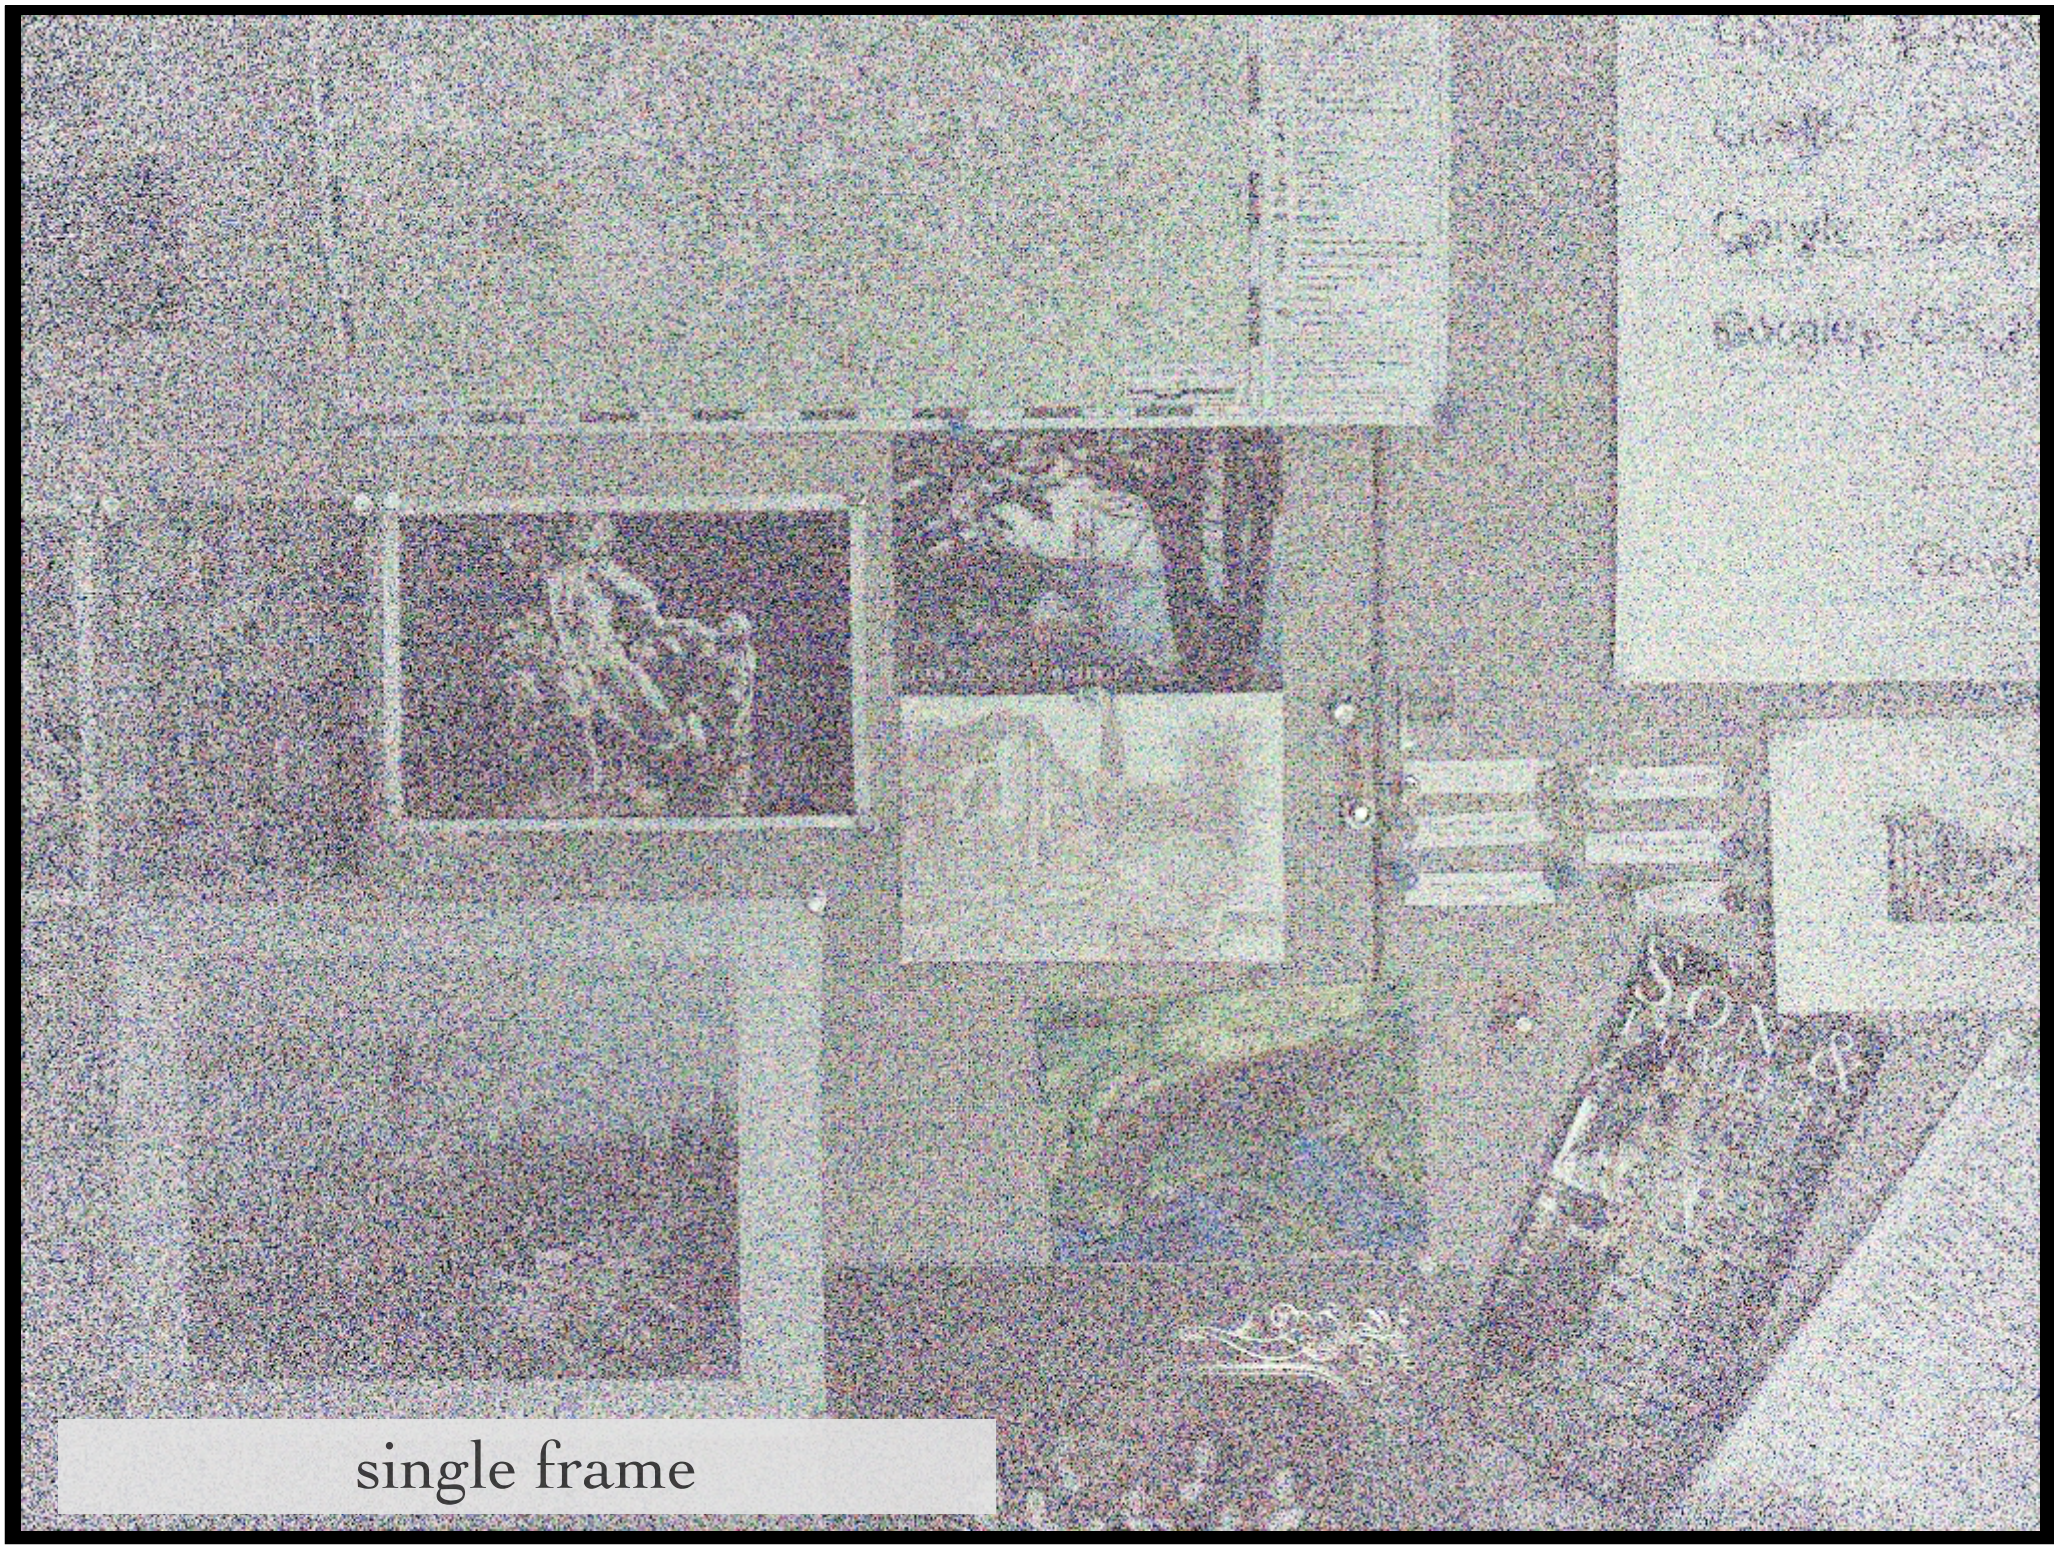

single frame

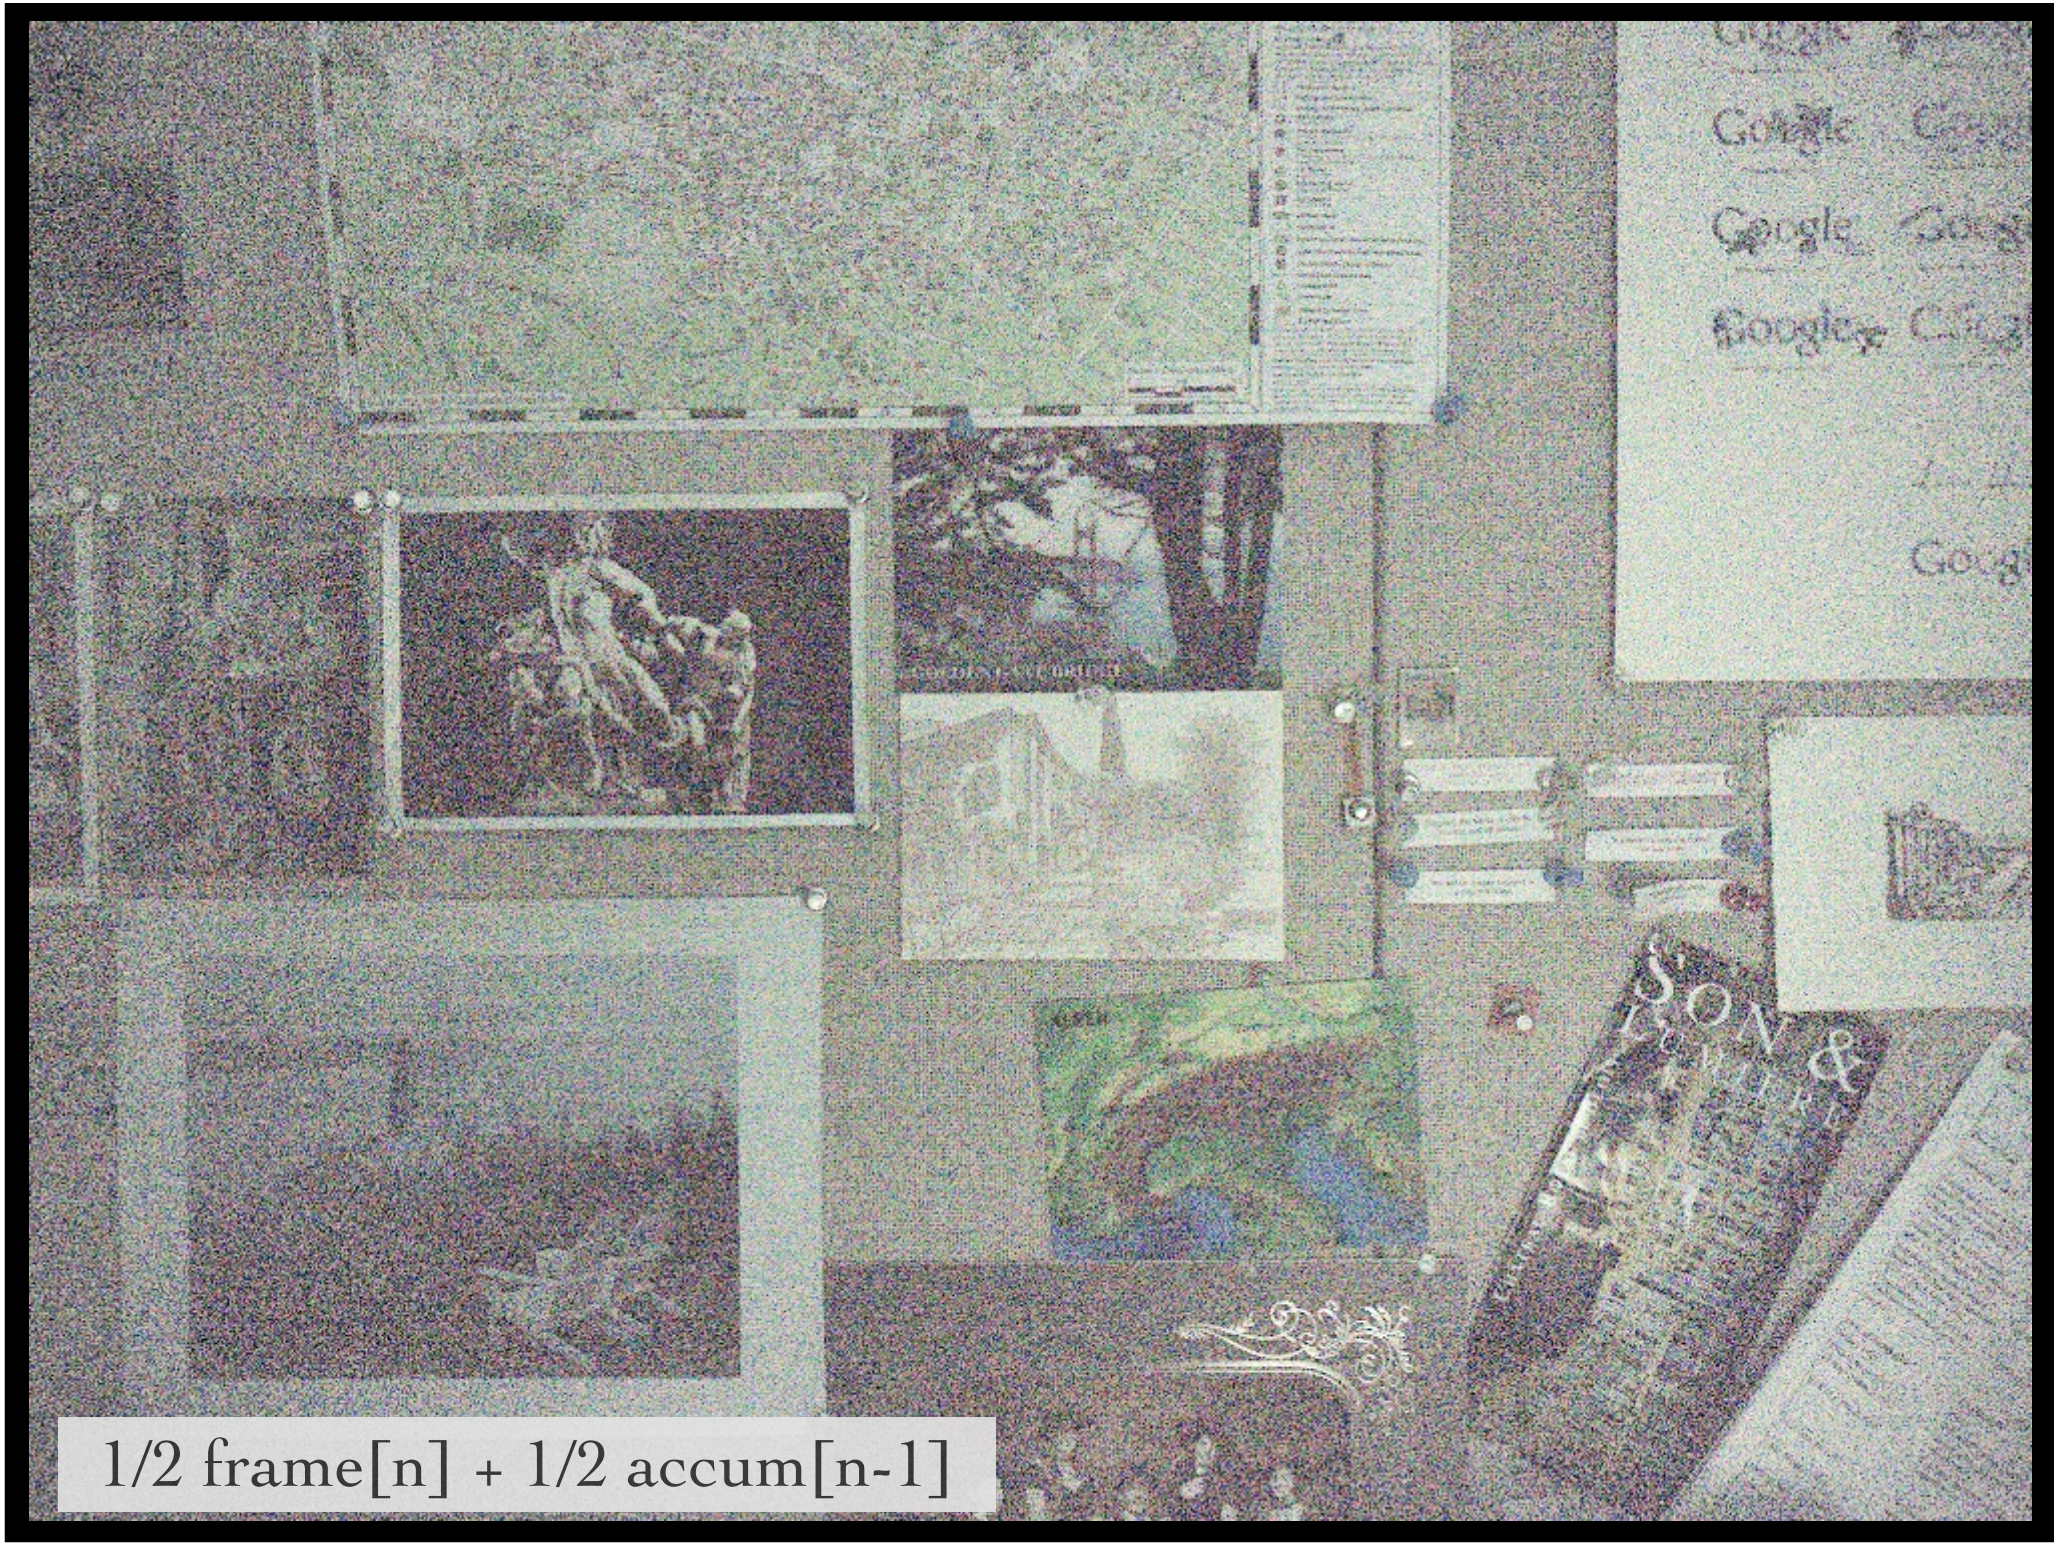
1/2 frame[n] + 1/2 accum[n-1]


1/4 frame[n] + 3/4 accum[n-1]


1/8 frame[n] + 7/8 accum[n-1]


GOLDEN GATE BRIDGE
ALPEN
1/16 frame[n] + 15/16 accum[n-1]

1/32 frame[n] + 31/32 accum[n-1]


GOLDEN GATE BRIDGE
ALPEN
1/64 frame[n] + 63/64 accum[n-1]

1/2 lux

1/64 frame[n] + 63/64 accum[n-1]

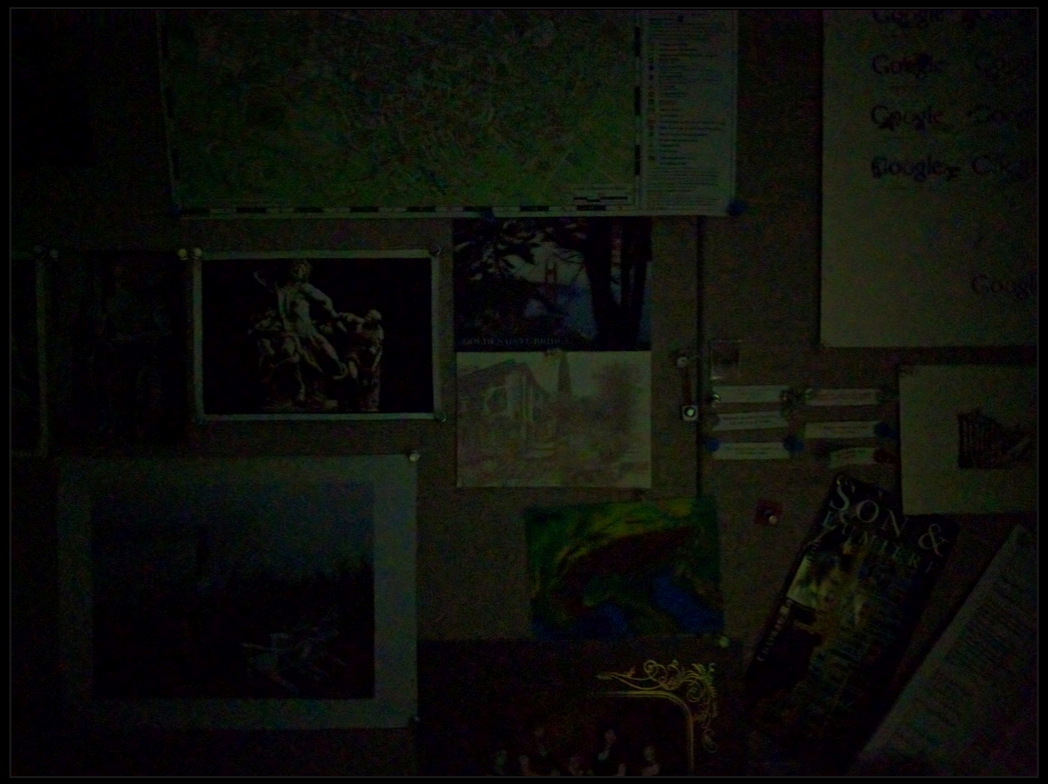

HDR+

HDR+

SeeInTheDark

ordinary cell phone picture

ordinary cell phone picture

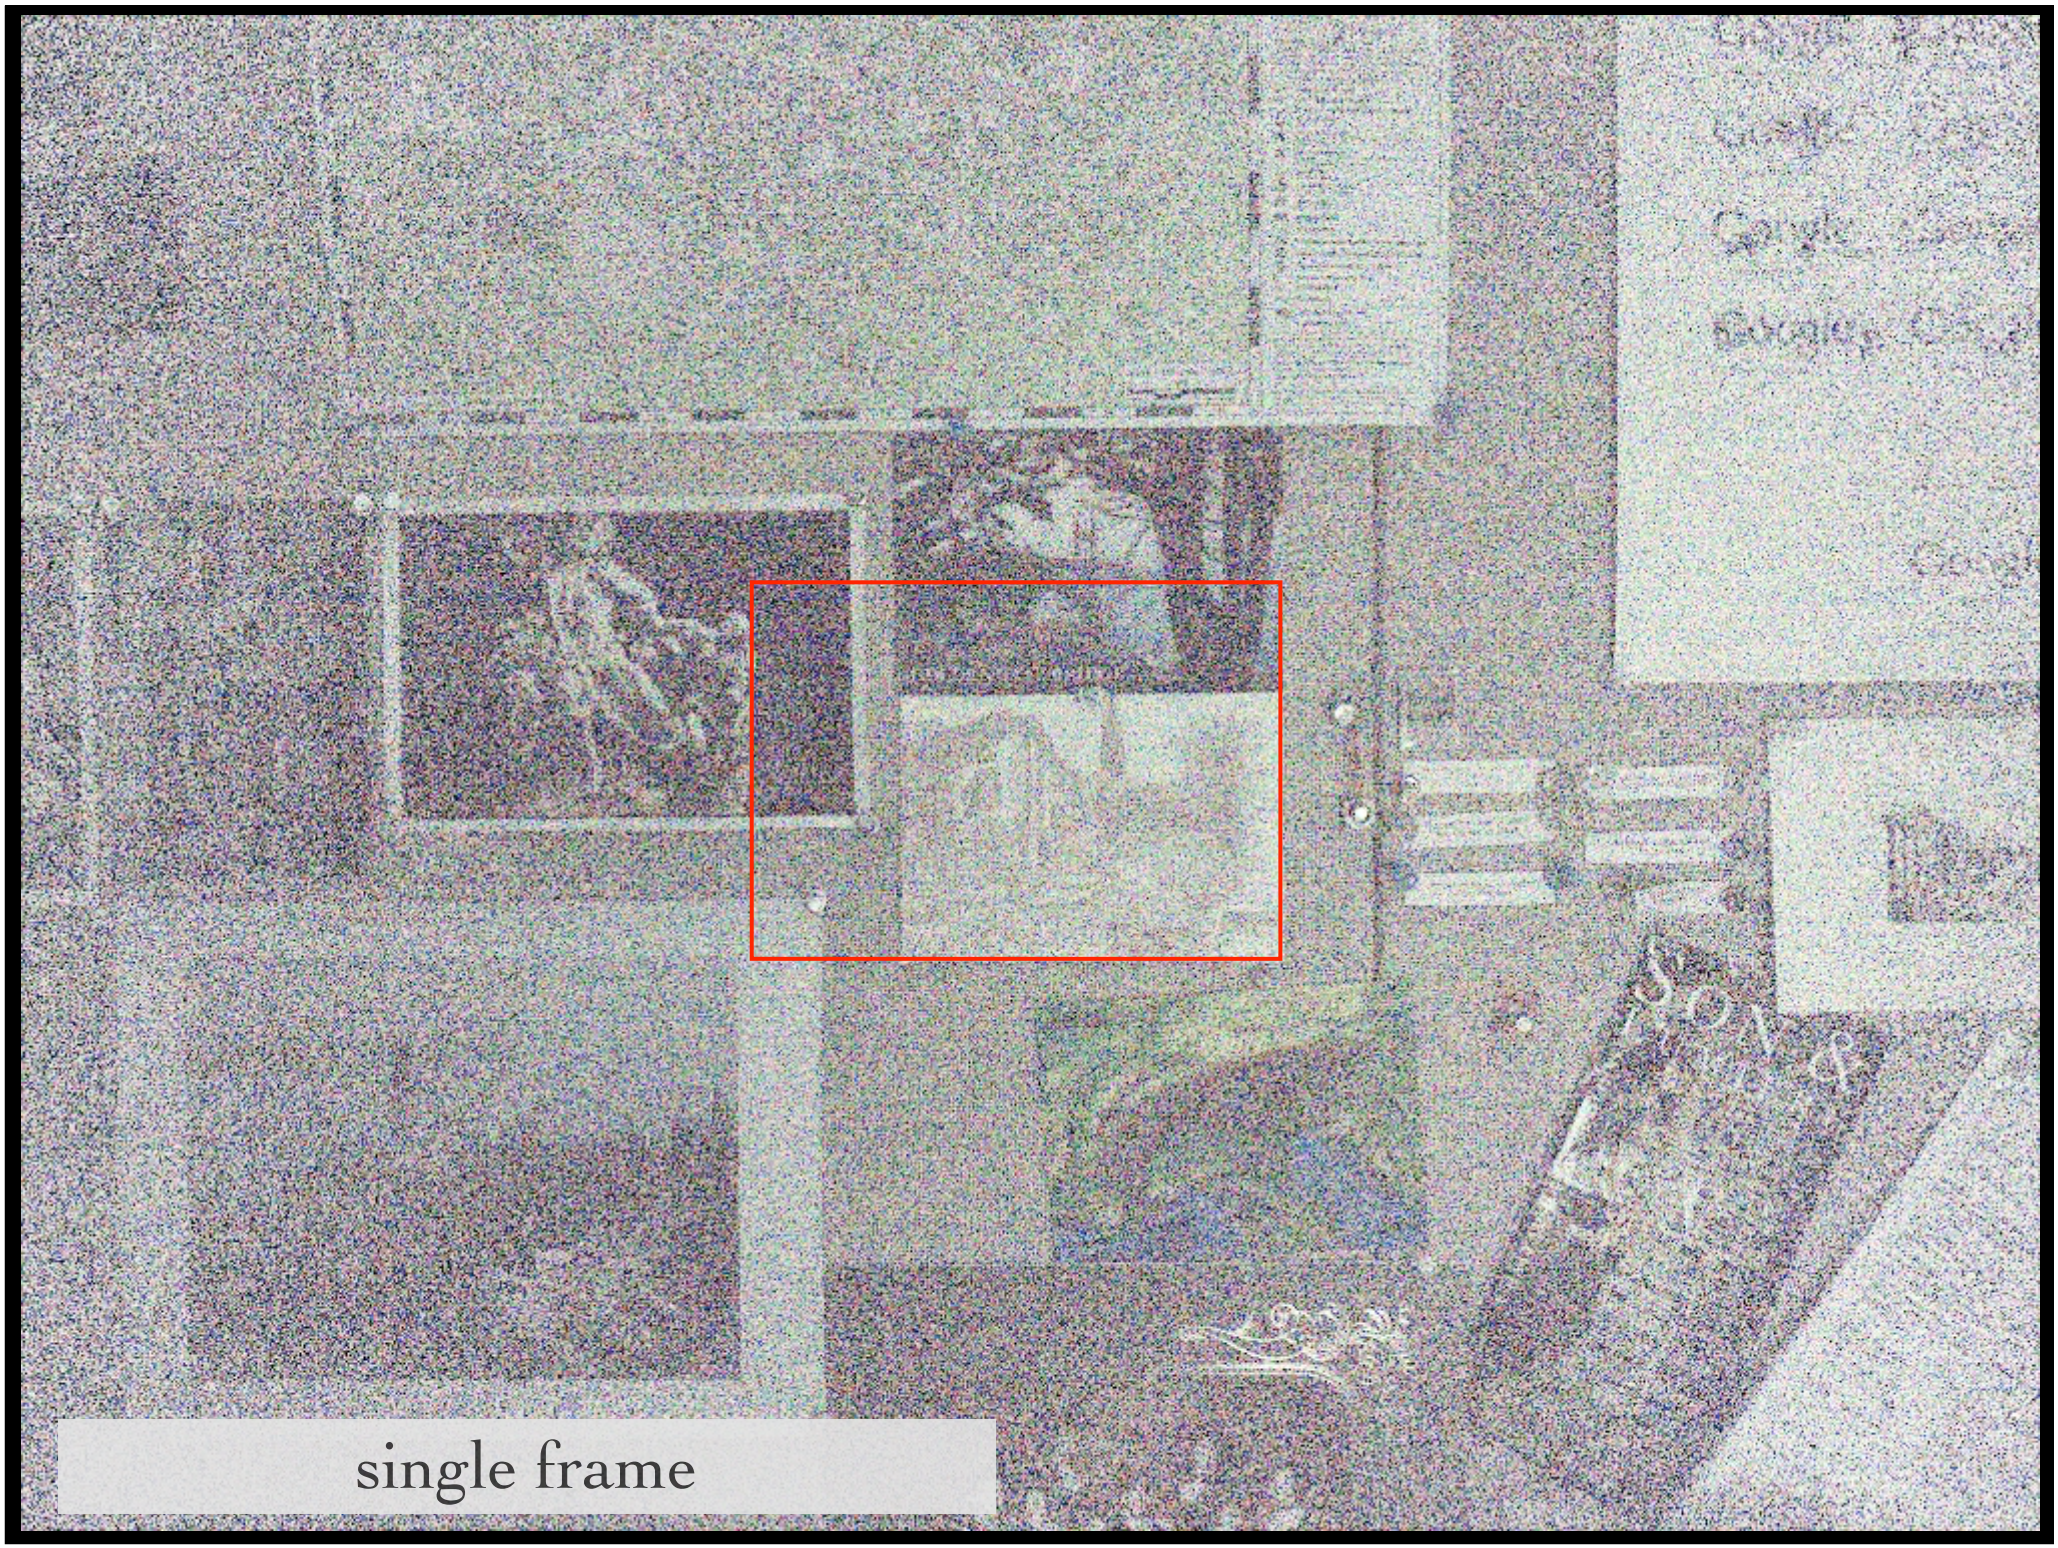
single frame


full-resolution single frame

WHAT IF THE SCENE
WERE DARKER?

AE/AF sequence (30fps)

AE/AF sequence (6fps)

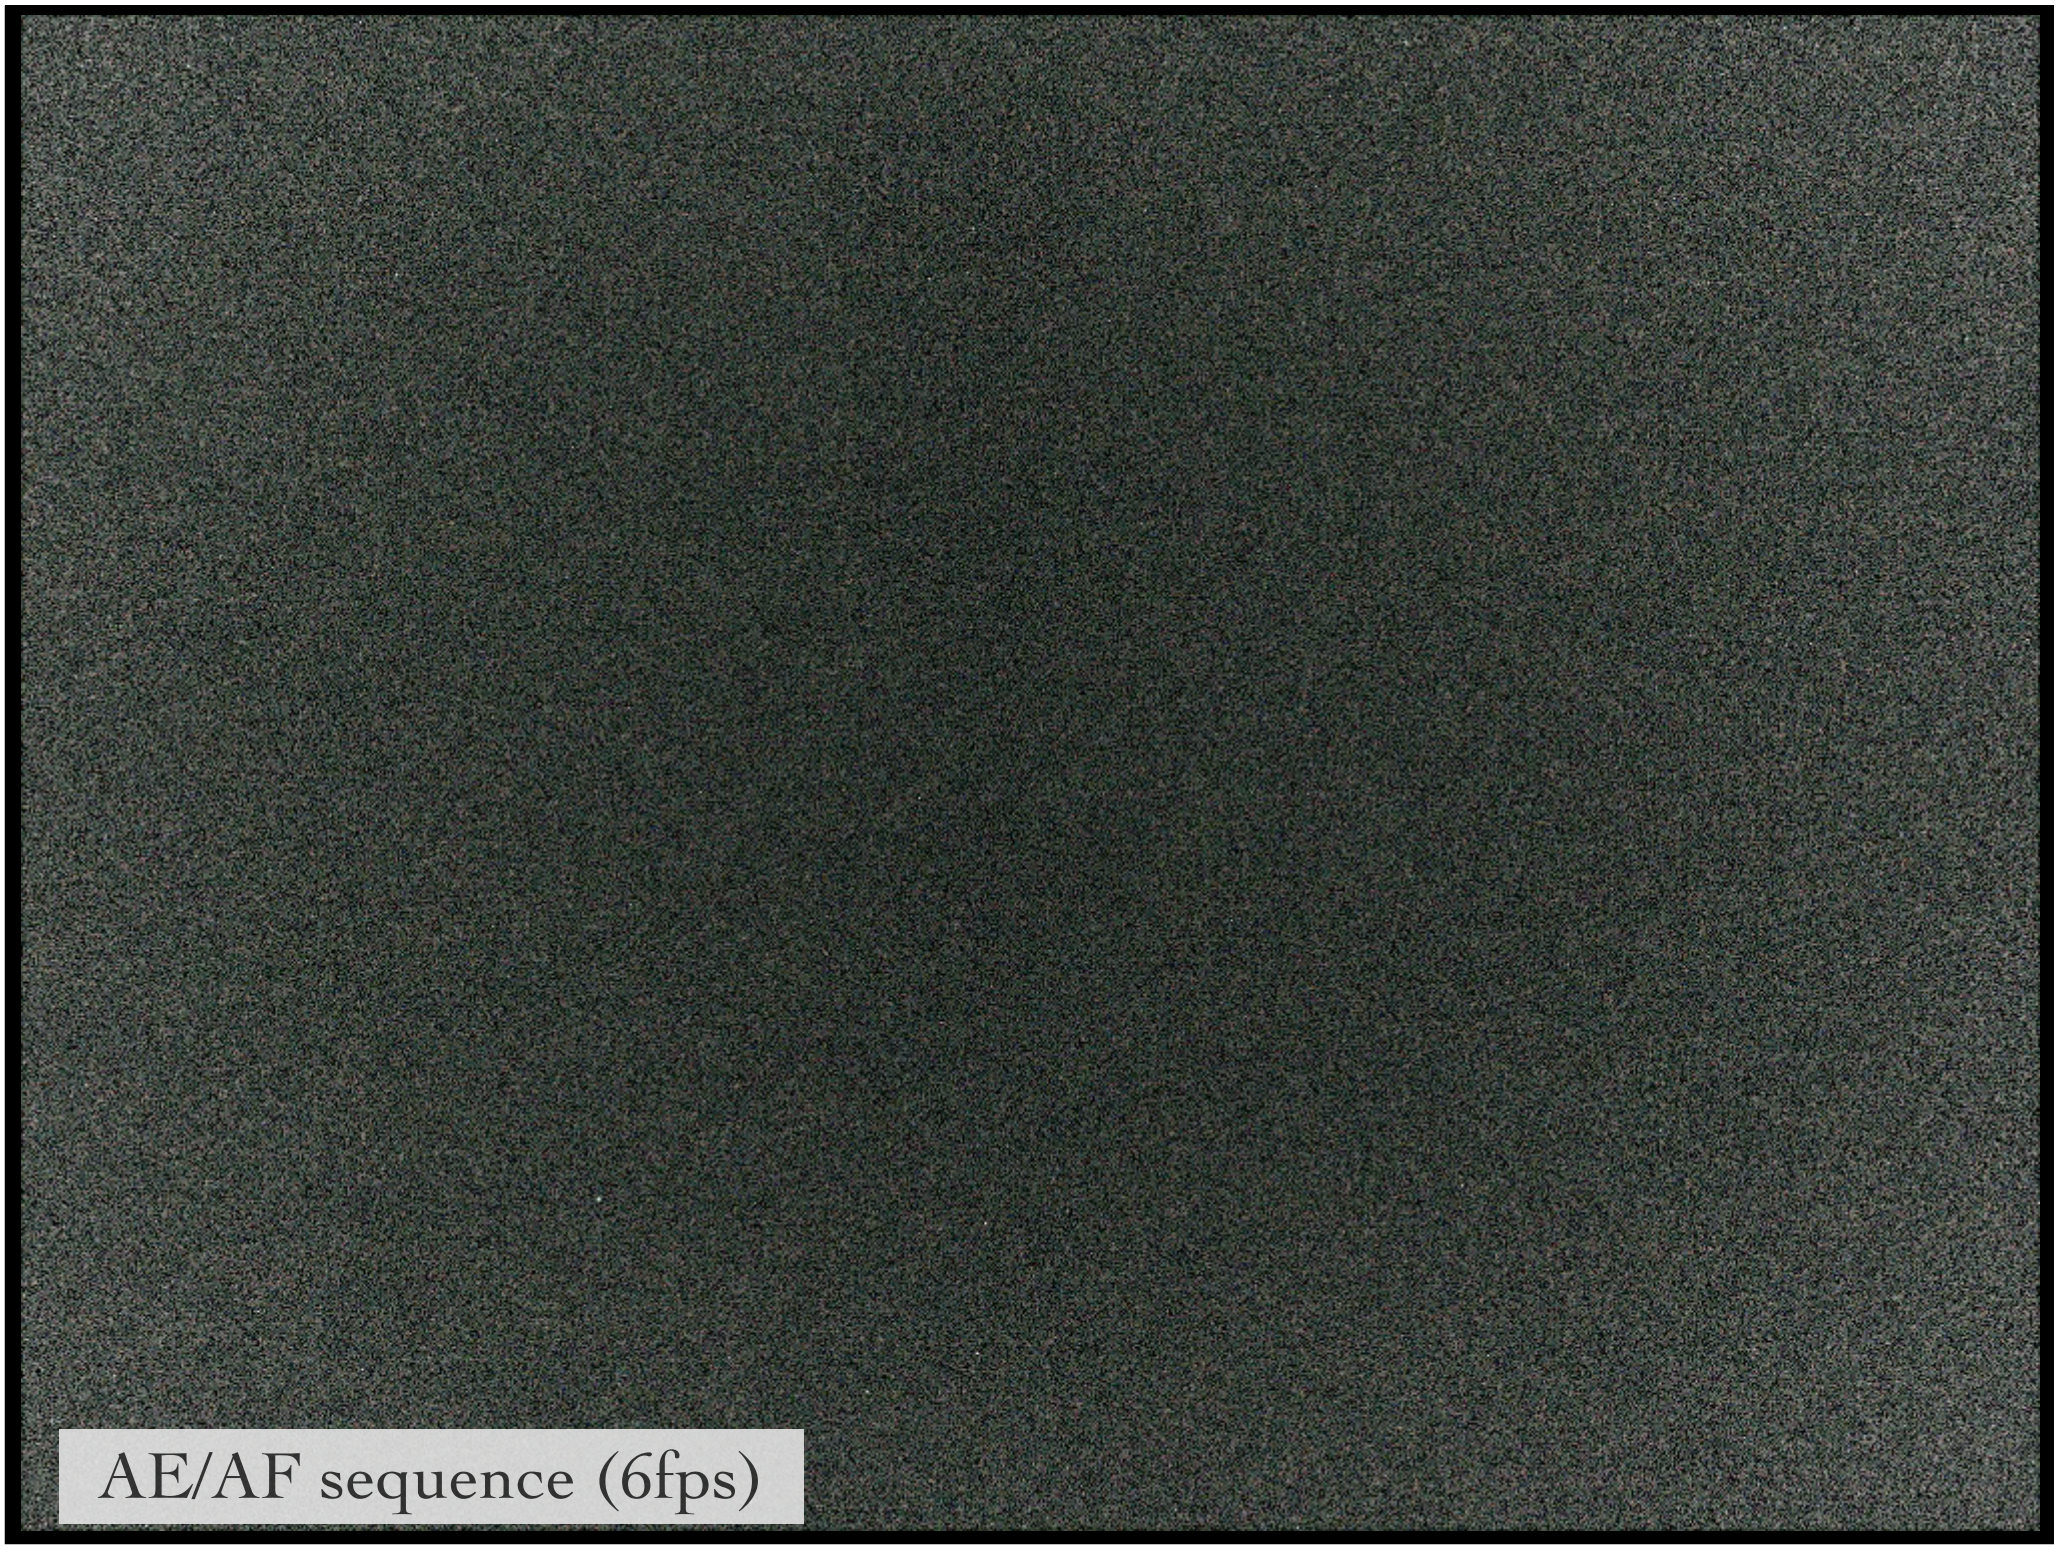
AE/AF sequence (6fps)

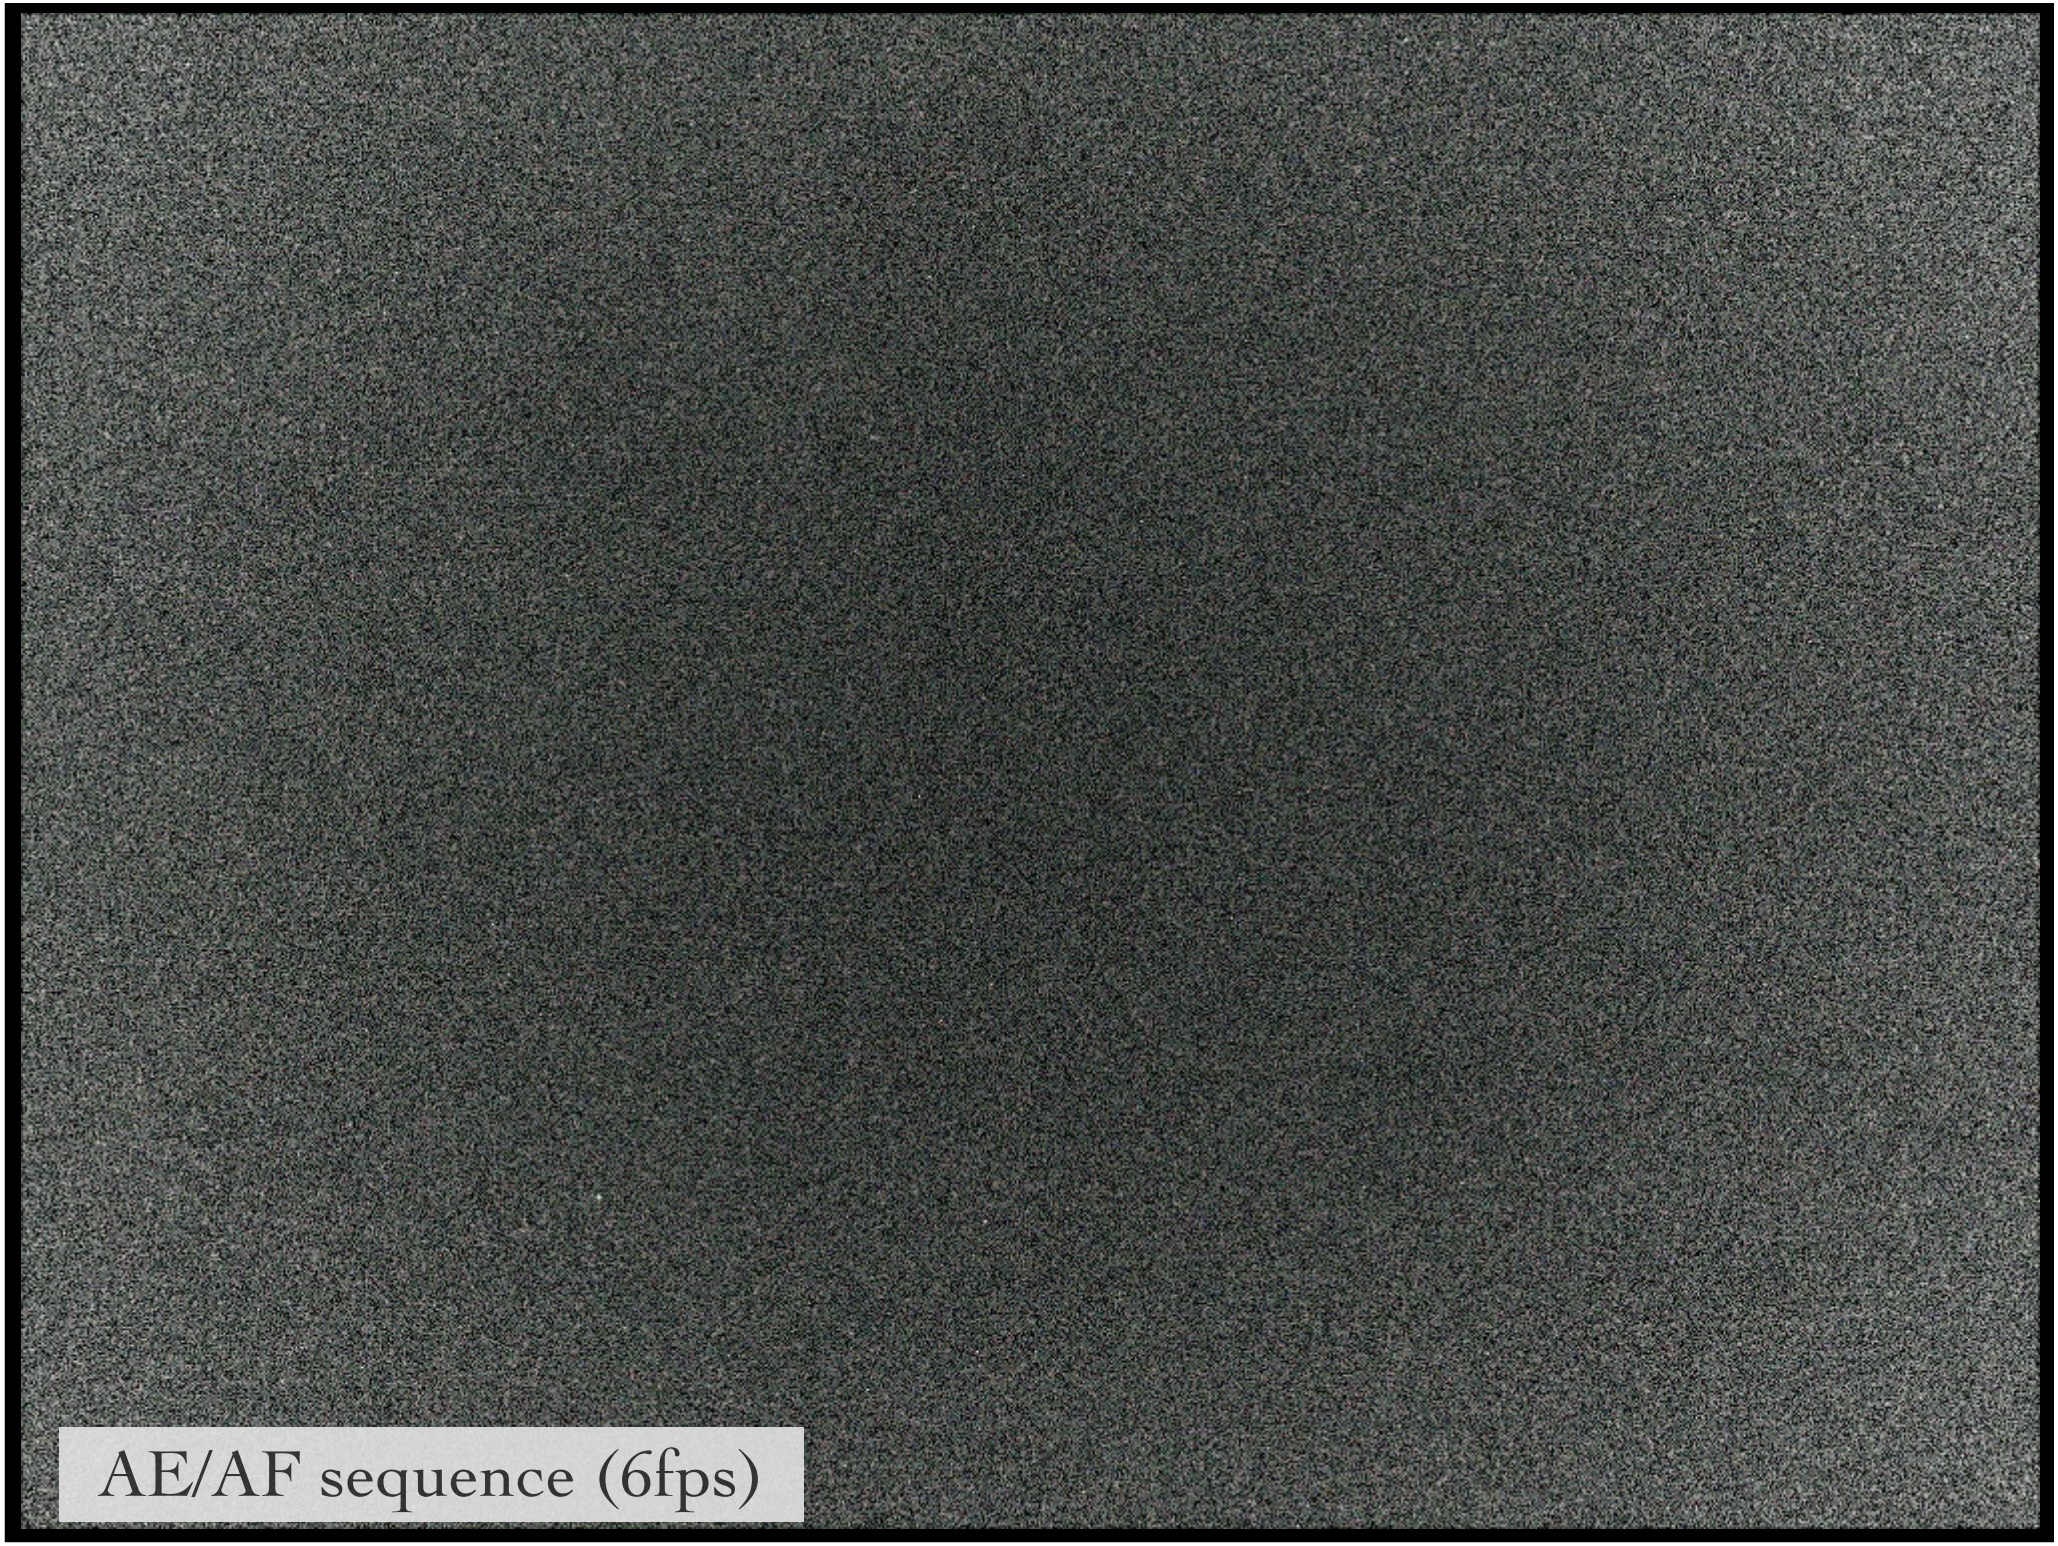
AE/AF sequence (6fps)


SKIL
DEWALT
0.1 lux
accumulating for 5 seconds

full moon, single frame

full moon, ~30 frames

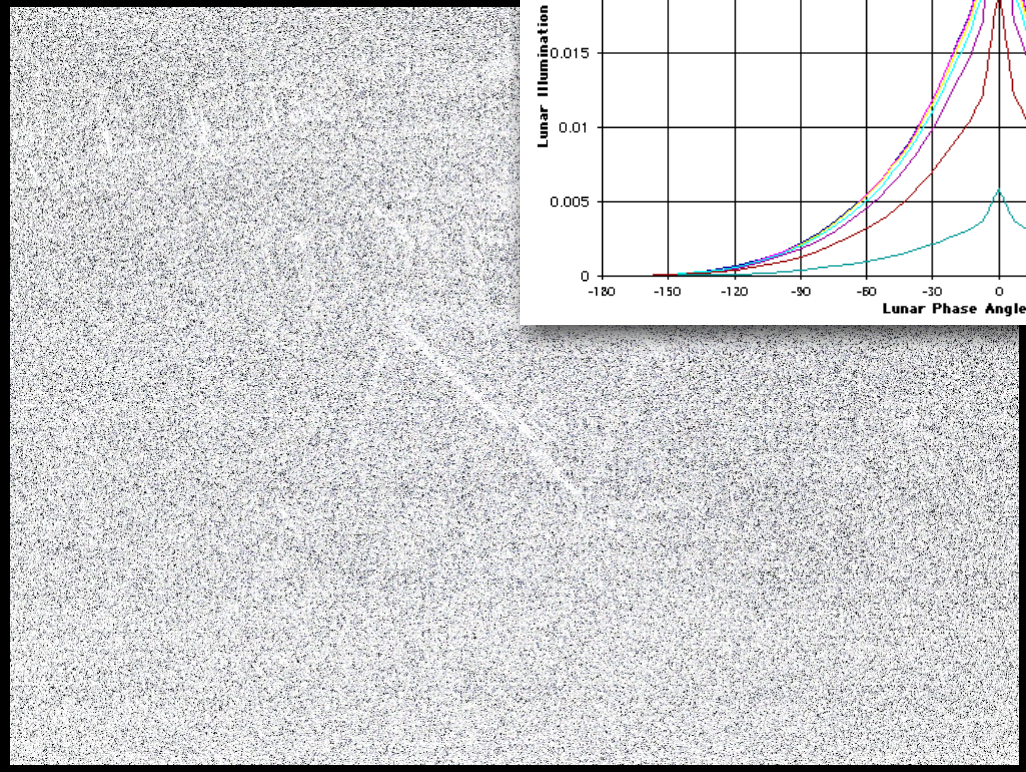

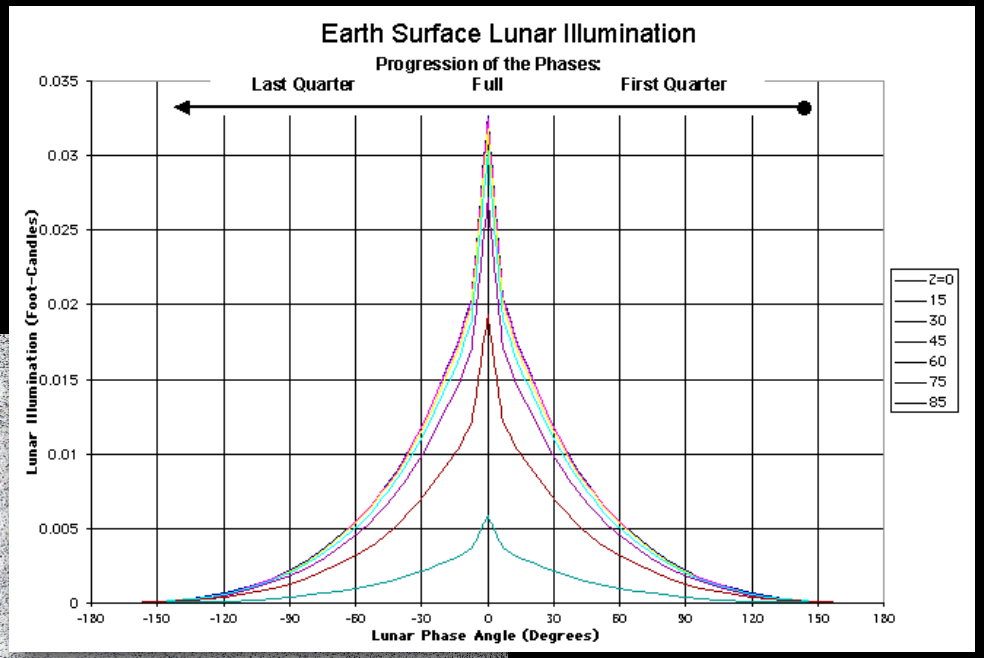

gibbous moon 30° above horizon

gibbous moon 20° above horizon

WHAT IF THE SCENE
WERE DARKER?

Signal-to-noise ratio

$$SNR = \frac{\text{mean pixel value}}{\text{standard deviation of pixel value}} = \frac{\mu}{\sigma}$$

$$= \frac{P\,Q_e\,t}{\sqrt{P\,Q_e\,t + D\,t + N_r^2}}$$

- where

P = incident photon flux (photons/pixel/sec)

Q_e = quantum efficiency

t = exposure time (sec)

D = dark current (electrons/pixel/sec), including hot pixels

N_r = read noise (rms electrons/pixel), including fixed pattern noise

(http://learn.hamamatsu.com/articles/ccdsnr.html)

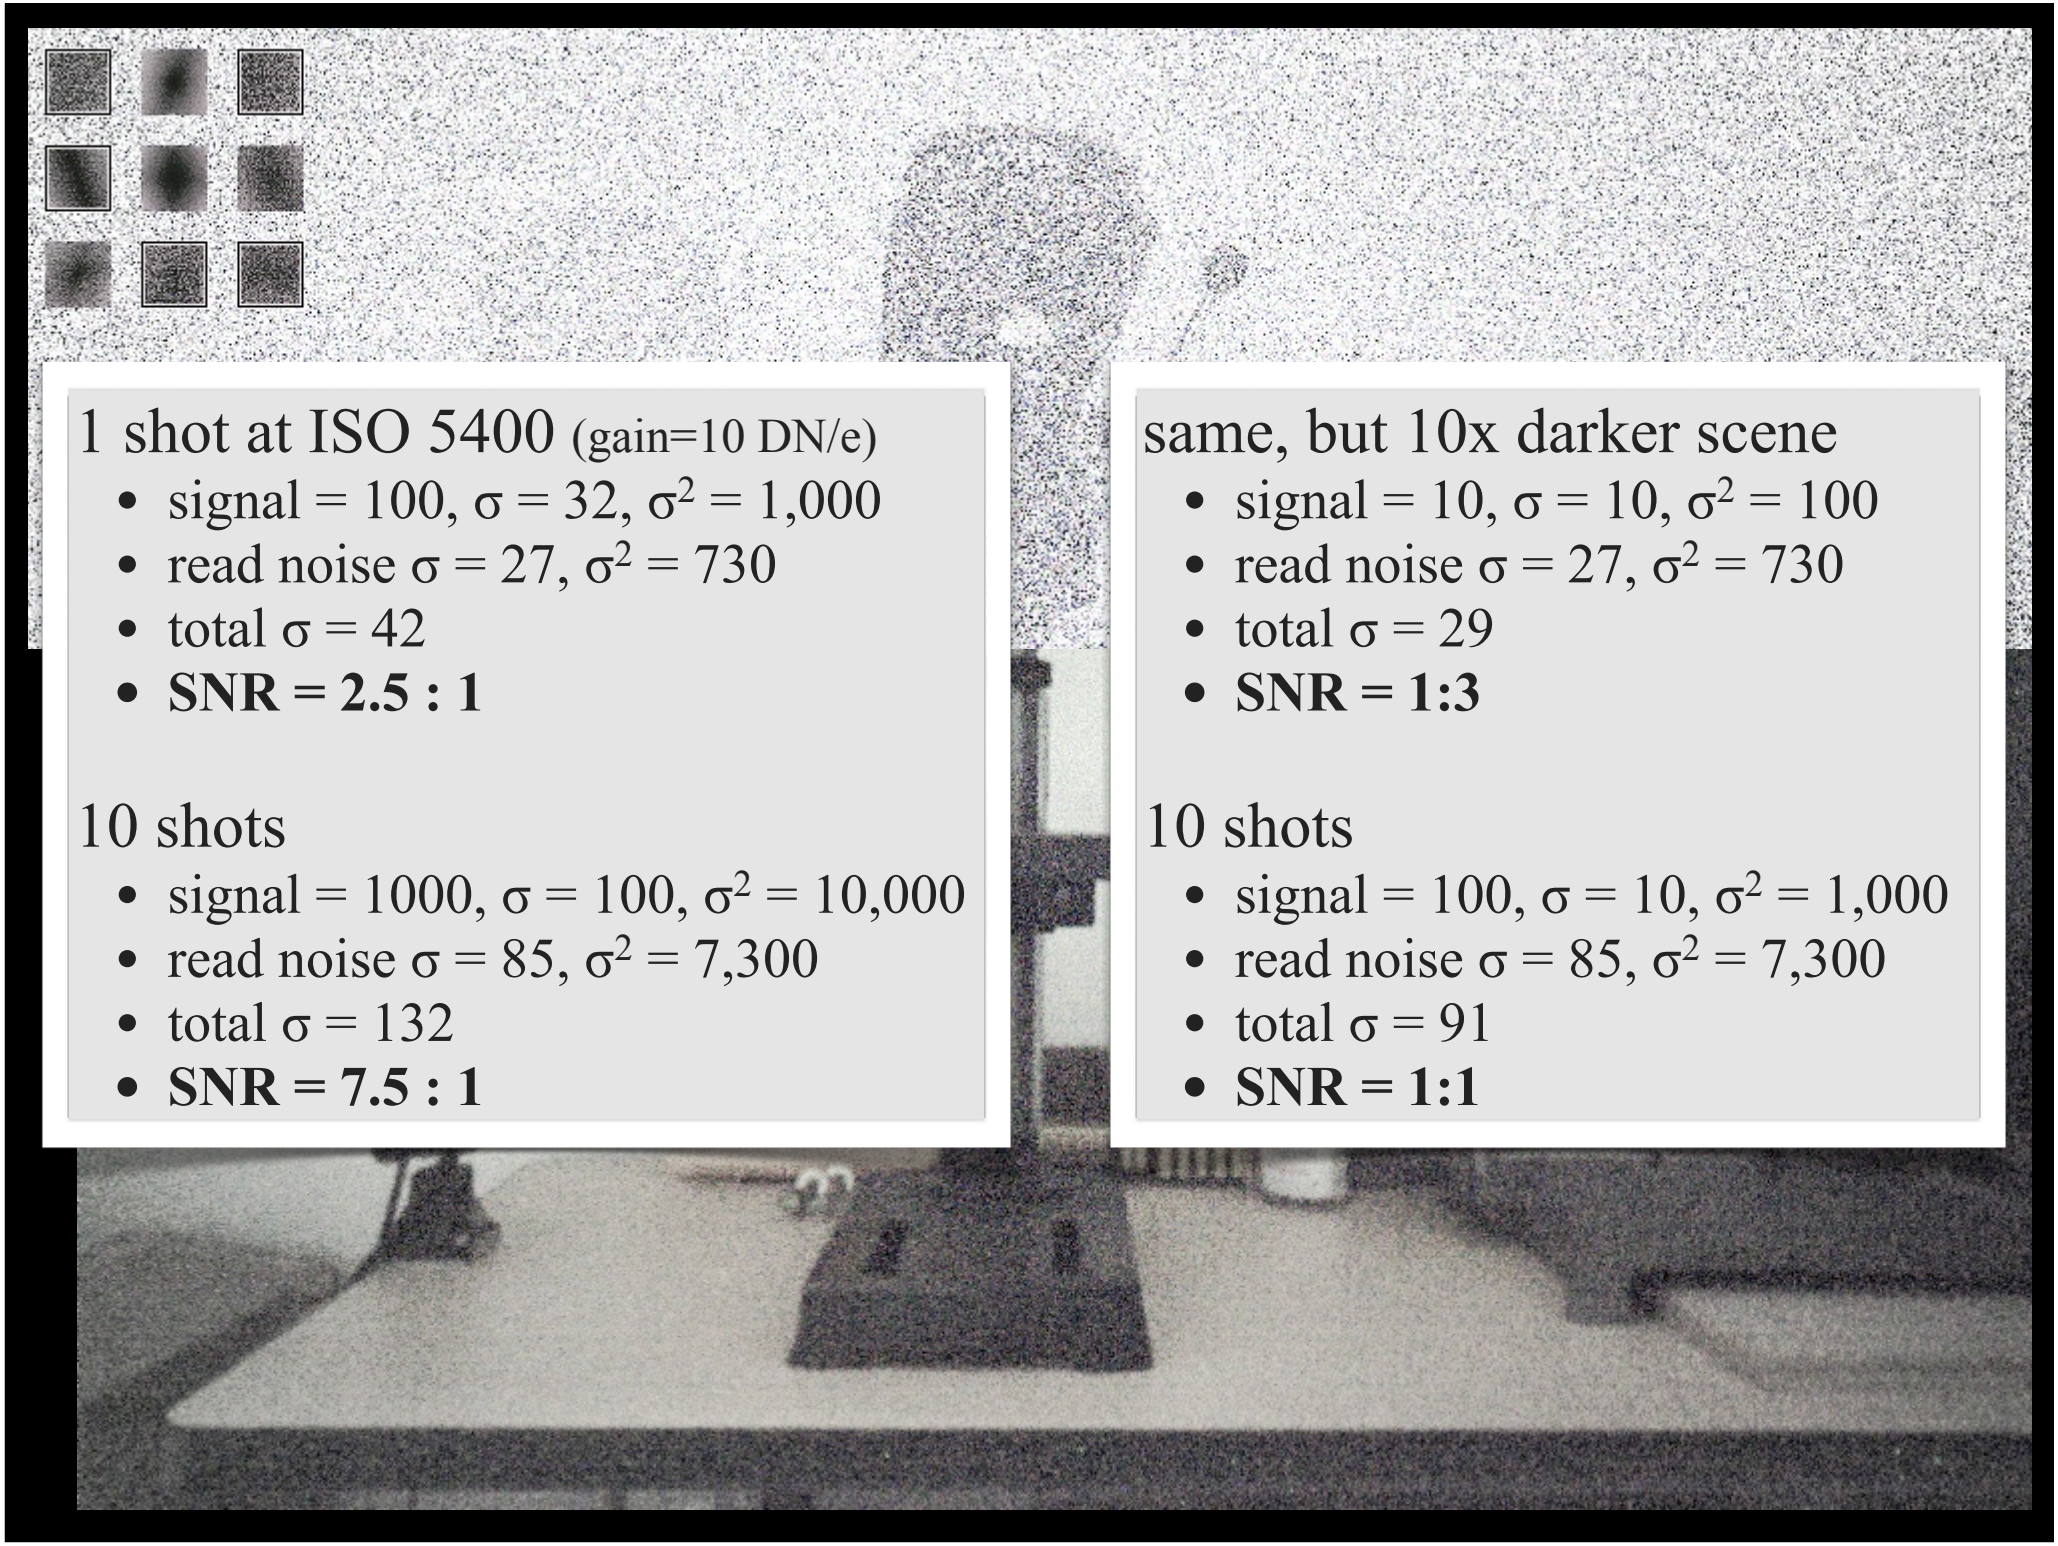

1 shot at ISO 5400 (gain=10 DN/e)

- signal = 100, $\sigma = 32$, $\sigma^2 = 1{,}000$
- read noise $\sigma = 27$, $\sigma^2 = 730$
- total $\sigma = 42$
- **SNR = 2.5 : 1**

10 shots

- signal = 1000, $\sigma = 100$, $\sigma^2 = 10{,}000$
- read noise $\sigma = 85$, $\sigma^2 = 7{,}300$
- total $\sigma = 132$
- **SNR = 7.5 : 1**

same, but 10x darker scene

- signal = 10, $\sigma = 10$, $\sigma^2 = 100$
- read noise $\sigma = 27$, $\sigma^2 = 730$
- total $\sigma = 29$
- **SNR = 1:3**

10 shots

- signal = 100, $\sigma = 10$, $\sigma^2 = 1{,}000$
- read noise $\sigma = 85$, $\sigma^2 = 7{,}300$
- total $\sigma = 91$
- **SNR = 1:1**

Where does the game stop?

- longer accumulation required to overcome read noise
- noise causes jitter in alignment, leading to blur in image
- not obvious how to auto-focus in very noisy images
- variation in noise sources over time
- scene doesn't hold still that long
- you can't hold still that long

- ways to cheat
  - IR illumination
  - more cameras...

Arrays of cameras

Stanford

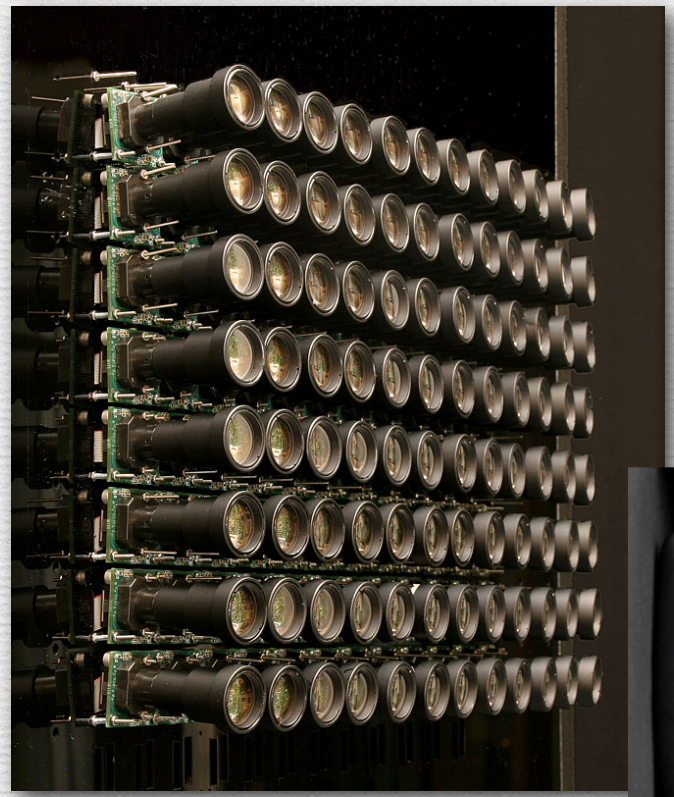

Light L16

fake

Light field camera arrays

Google Jump

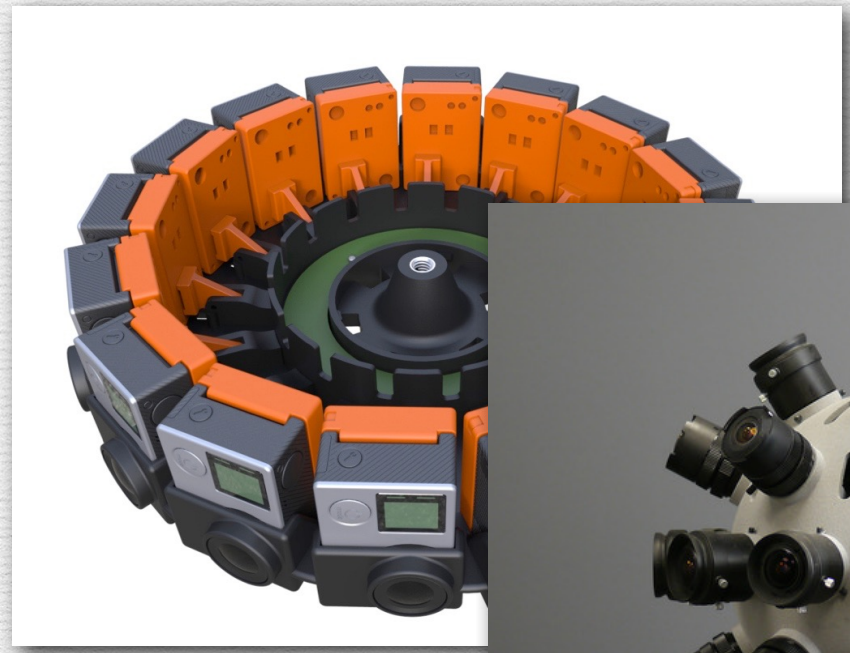

Jaunt

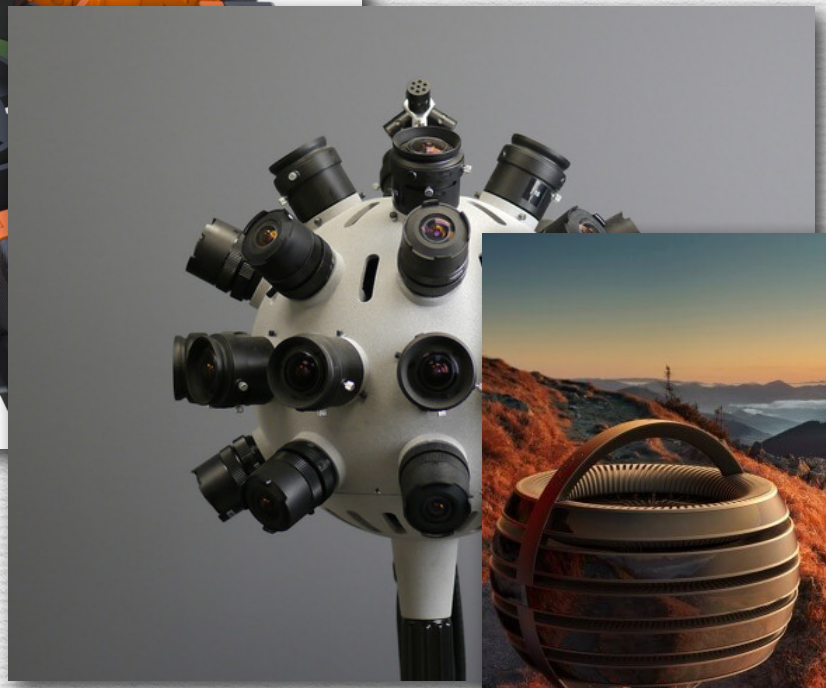

Lytro

(Eddy Talvala, Nexus 6P, 240fps, 2015, available on YouTube at https://www.youtube.com/watch?v=UaG8DuvpOLU)

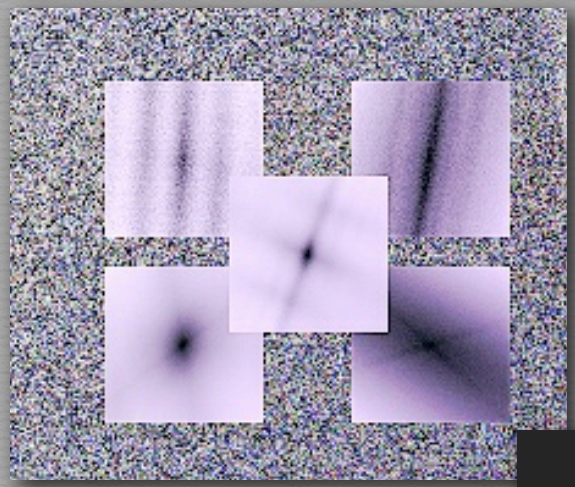


```
Sensor sensor;
Shot low, med, high;

low.exposure = 1/80.;
med.exposure = 1/20.;
high.exposure = 1/5.;
```